

# Critical Reasoning: A User's Manual

© Chris Swoyer

Version 3.0 (5/30/2002)



# Contents

<b>I</b>	<b>Basic Concepts of Critical Reasoning</b>	<b>1</b>
<b>1</b>	<b>Basic Concepts of Critical Reasoning</b>	<b>5</b>
1.1	Basic Concepts . . . . .	5
1.2	A Role for Reason . . . . .	13
1.3	Improving Reasoning . . . . .	17
1.4	Chapter Exercises . . . . .	18
<b>II</b>	<b>Reasons and Arguments</b>	<b>19</b>
<b>2</b>	<b>Arguments</b>	<b>23</b>
2.1	Arguments . . . . .	24
2.1.1	Inferences and Arguments . . . . .	24
2.2	Uses of Arguments . . . . .	25
2.2.1	Reasoning . . . . .	25
2.2.2	Persuasion . . . . .	25
2.2.3	Evaluation . . . . .	25
2.3	Identifying Arguments in their Natural Habitat . . . . .	26
2.3.1	Indicator Words . . . . .	26
2.4	Putting Arguments into Standard Form . . . . .	27
2.4.1	Arguments vs. Conditionals . . . . .	27
2.5	Deductive Validity . . . . .	30
2.5.1	Definition of Deductive Validity . . . . .	30
2.5.2	Further Features of Deductive Validity . . . . .	33
2.5.3	Soundness . . . . .	33
2.6	Method of Counterexample . . . . .	34
2.7	Inductive Strength . . . . .	39
2.8	Evaluating Arguments . . . . .	41
2.9	Chapter Exercises . . . . .	42

<b>3</b>	<b>Conditionals and Conditional Arguments</b>	<b>43</b>
3.1	Conditionals and their Parts . . . . .	43
3.1.1	Alternative Ways to State Conditionals . . . . .	45
3.2	Necessary and Sufficient Conditions . . . . .	48
3.3	Conditional Arguments . . . . .	53
3.3.1	Conditional Arguments that Affirm . . . . .	53
3.3.2	Conditional Arguments that Deny . . . . .	54
3.4	Chapter Exercises . . . . .	55
<b>III</b>	<b>The Acquisition and Retention of Information</b>	<b>61</b>
<b>4</b>	<b>Perception: Expectation and Inference</b>	<b>65</b>
4.1	Perception and Reasoning . . . . .	66
4.2	Perception is Selective . . . . .	66
4.3	There's More to Seeing than Meets the Eye . . . . .	67
4.3.1	Information Processing . . . . .	68
4.4	Going Beyond the Information Given . . . . .	68
4.5	Perception and Inference . . . . .	70
4.6	What Ambiguous Figures Teach Us . . . . .	71
4.7	Perceptual Set: the Role of Expectations . . . . .	73
4.7.1	Classification and Set . . . . .	75
4.7.2	Real-life Examples . . . . .	76
4.8	There's more to Hearing, Feeling, ... . . . .	77
4.8.1	Hearing . . . . .	77
4.8.2	Feelings . . . . .	78
4.9	Seeing What We Want to See . . . . .	79
4.9.1	Perception as Inference . . . . .	81
4.9.2	Seeing Shouldn't be Believing . . . . .	81
4.10	Chapter Exercises . . . . .	83
<b>5</b>	<b>Evaluating Sources of Information</b>	<b>85</b>
5.1	Other People as Sources of Information . . . . .	86
5.1.1	Information: We need something to Reason About . . . . .	86
5.2	Expertise . . . . .	88
5.2.1	What is an Expert? . . . . .	88
5.2.2	Fields of Expertise . . . . .	89
5.3	Evaluating Claims to Expertise . . . . .	89
5.4	Who Do we Listen To? . . . . .	93
5.4.1	Faking Expertise: The Aura of Authority . . . . .	94

---

5.4.2	Appearing to Go Against Self-interest . . . . .	96
5.5	Evaluating Testimony in General . . . . .	97
5.6	Safeguards . . . . .	99
5.7	Chapter Exercises . . . . .	102
<b>6</b>	<b>The Net: Finding and Evaluating Information on the Web</b>	<b>107</b>
6.1	The World Wide Web . . . . .	108
6.1.1	What is the World Wide Web? . . . . .	108
6.1.2	Bookmarks . . . . .	109
6.2	Search Engines . . . . .	110
6.2.1	Specific Search Engines . . . . .	111
6.2.2	Metasearch Engines . . . . .	112
6.2.3	Specialty Search Engines . . . . .	113
6.2.4	Rankings of Results . . . . .	113
6.3	Refining your Search . . . . .	113
6.4	Evaluating Material on the Net . . . . .	115
6.4.1	Stealth Advocacy . . . . .	119
6.5	Evaluation Checklist . . . . .	120
6.6	Citing Information from the Net . . . . .	121
6.7	Chapter Exercises . . . . .	122
<b>7</b>	<b>Memory and Reasoning</b>	<b>125</b>
7.1	Memory and Reasoning . . . . .	126
7.2	Stages in Memory . . . . .	126
7.2.1	Where Things can go Wrong . . . . .	128
7.3	Encoding . . . . .	128
7.4	Storage . . . . .	129
7.4.1	Editing and Revising . . . . .	129
7.5	Retrieval . . . . .	131
7.5.1	Context and Retrieval Cues . . . . .	134
7.5.2	Schemas . . . . .	135
7.6	Summary: Inference and Influences on Memory . . . . .	137
7.7	Chapter Exercises . . . . .	137
<b>8</b>	<b>Memory II: Pitfalls and Remedies</b>	<b>139</b>
8.1	Misattribution of Source . . . . .	140
8.2	The Power of Suggestion and the Misinformation Effect . . . . .	140
8.3	Confidence and Accuracy . . . . .	141
8.3.1	Flashbulb Memories . . . . .	141
8.4	False Memories . . . . .	142

8.4.1	Motivated Misremembering . . . . .	143
8.4.2	Childhood Trauma and False-Memory Syndrome . . . . .	143
8.5	Belief Perseveration . . . . .	144
8.6	Hindsight Bias . . . . .	145
8.7	Inert Knowledge . . . . .	145
8.8	Eyewitness Testimony . . . . .	146
8.9	Primacy and Recency Effects . . . . .	147
8.9.1	The Primacy Effect . . . . .	147
8.9.2	The Recency Effect . . . . .	147
8.10	Collective Memory . . . . .	148
8.11	Remedies . . . . .	149
8.11.1	Safeguards . . . . .	149
8.11.2	Ways to Improve Memory . . . . .	149
8.12	Chapter Summary . . . . .	151
8.13	Chapter Exercises . . . . .	152
8.14	Appendix: Different Memory Systems and False Memories . . . . .	154
<b>9</b>	<b>Emotions and Reasoning</b>	<b>157</b>
9.1	The Pervasiveness of Emotions . . . . .	158
9.1.1	Emotions and Information . . . . .	158
9.2	Stress . . . . .	159
9.3	Legitimate Appeals to Emotion . . . . .	160
9.4	Illegitimate Appeals to Emotion . . . . .	161
9.4.1	Pity . . . . .	161
9.4.2	Fear . . . . .	161
9.4.3	Anger . . . . .	162
9.5	Self-serving Biases . . . . .	163
9.6	Chapter Exercises . . . . .	165
<b>IV</b>	<b>Relevance, Irrelevance, and Fallacies</b>	<b>169</b>
<b>10</b>	<b>Relevance, Irrelevance, and Reasoning</b>	<b>173</b>
10.1	Relevance . . . . .	174
10.2	Fallacy of Irrelevant Reasons . . . . .	178
10.3	Arguments Against the Person . . . . .	180
10.4	The Strawman Fallacy . . . . .	183
10.4.1	Safeguards . . . . .	185
10.5	Appeal to Ignorance . . . . .	187
10.5.1	Burden of Proof . . . . .	188

---

10.6 Suppressed (or Neglected) Evidence . . . . .	191
10.7 Chapter Exercises . . . . .	192
<b>11 Fallacies: Common Ways of Reasoning Badly</b>	<b>195</b>
11.1 Begging the Question . . . . .	196
11.2 The Either/Or Fallacy . . . . .	203
11.2.1 Clashes of Values . . . . .	206
11.2.2 Safeguards . . . . .	208
11.3 Drawing the Line . . . . .	212
11.4 Inconsistency . . . . .	213
11.5 Chapter Exercises . . . . .	214
11.6 Summary of Fallacies . . . . .	217
<b>V Induction and Probability</b>	<b>219</b>
<b>12 Induction in the Real World</b>	<b>223</b>
12.1 Life is Uncertain . . . . .	223
12.2 Inductively Strong Arguments . . . . .	226
12.3 Chapter Exercises . . . . .	227
<b>13 Rules for Calculating Probabilities</b>	<b>229</b>
13.1 Intuitive Illustrations . . . . .	230
13.2 Probabilities are Numbers . . . . .	230
13.2.1 Notation . . . . .	231
13.3 Rules for Calculating Probabilities . . . . .	233
13.3.1 Absolutely Certain Outcomes . . . . .	233
13.3.2 Negations . . . . .	234
13.3.3 Disjunctions with Incompatible Disjuncts . . . . .	236
13.4 More Rules for Calculating Probabilities . . . . .	240
13.4.1 Conjunctions with Independent Conjuncts . . . . .	240
13.4.2 Disjunctions with Compatible Disjuncts . . . . .	243
13.5 Chapter Exercises . . . . .	244
13.6 Appendix: Working with Fractions . . . . .	246
<b>14 Conditional Probabilities</b>	<b>249</b>
14.1 Conditional Probabilities . . . . .	249
14.1.1 Characterization of Conditional Probability . . . . .	251
14.1.2 The General Conjunction Rule . . . . .	253
14.2 Analyzing Probability Problems . . . . .	254

14.2.1	Examples of Problem Analysis . . . . .	255
14.3	Odds and Ends . . . . .	259
14.3.1	Sample Problems with Answers . . . . .	262
14.3.2	More Complex Problems . . . . .	263
14.4	Chapter Exercises . . . . .	264
14.4.1	Summary of Rules . . . . .	267
<b>VI</b>	<b>Induction in the Real World</b>	<b>269</b>
<b>15</b>	<b>Samples and Correlations</b>	<b>273</b>
15.1	Descriptive Statistics . . . . .	274
15.1.1	Features of Samples . . . . .	275
15.1.2	Exercises . . . . .	277
15.2	Inferences from Samples to Populations . . . . .	278
15.2.1	Sampling in Every day Life . . . . .	278
15.2.2	Samples and Inference . . . . .	279
15.2.3	Good Samples . . . . .	279
15.2.4	Bad Sampling and Bad Reasoning . . . . .	283
15.2.5	Exercises . . . . .	284
15.3	Correlation . . . . .	287
15.3.1	Correlation is Comparative . . . . .	289
15.3.2	Exercises . . . . .	294
15.4	Real vs. Illusory Correlations . . . . .	295
15.4.1	Ferretting out Illusory Correlations . . . . .	297
15.4.2	The Halo Effect: A Case Study in Illusory Correlation . . . . .	299
15.5	Chapter Exercises . . . . .	301
<b>16</b>	<b>Applications and Pitfalls</b>	<b>303</b>
16.1	What do the Numbers Mean? . . . . .	304
16.1.1	Ratios of Successes to Failures . . . . .	304
16.1.2	Frequencies . . . . .	305
16.1.3	Degrees of Belief . . . . .	305
16.1.4	How can we Comprehend such Tiny Numbers? . . . . .	306
16.1.5	Probabilistic Reasoning without Numbers . . . . .	307
16.2	Expected Value . . . . .	308
16.2.1	Pascal's Wager . . . . .	310
16.3	The Gambler's Fallacy . . . . .	311
16.4	The Conjunction Fallacy . . . . .	312
16.5	Doing Better by Using Frequencies . . . . .	314



---

16.6	Why Things go Wrong . . . . .	315
16.7	Regression to the Mean . . . . .	317
16.7.1	Regression and Reasoning . . . . .	319
16.8	Coincidence . . . . .	320
16.9	Chapter Exercises . . . . .	321
<b>VII</b>	<b>Systematic Biases and Distortions in Reasoning</b>	<b>325</b>
<b>17</b>	<b>Heuristics and Biases</b>	<b>329</b>
17.1	Inferential Heuristics . . . . .	330
17.1.1	Sampling Revisited . . . . .	330
17.2	The Availability Heuristic . . . . .	331
17.2.1	Why Things Are Available . . . . .	332
17.3	The Representativeness Heuristic . . . . .	335
17.3.1	Specificity Revisited . . . . .	337
17.4	Base-Rates . . . . .	337
17.5	Anchoring and Adjustment . . . . .	340
17.5.1	Anchoring Effects can be Very Strong . . . . .	341
17.5.2	Anchoring and Adjustment in the Real World . . . . .	341
17.5.3	Safeguards . . . . .	342
17.6	Chapter Exercises . . . . .	342
<b>18</b>	<b>More Biases, Pitfalls, and Traps</b>	<b>347</b>
18.1	Framing Effects . . . . .	348
18.1.1	Different Presentations of Alternatives . . . . .	348
18.1.2	Losses vs. Gains . . . . .	349
18.1.3	Loss Aversion . . . . .	351
18.1.4	The Certainty Effect . . . . .	352
18.2	Psychological Accounting . . . . .	352
18.3	Magic Numbers . . . . .	353
18.4	Sunk Costs . . . . .	358
18.5	Confirmation Bias . . . . .	359
18.6	Self-Fulfilling Prophecies . . . . .	361
18.7	The Validity Effect and Mere Exposure . . . . .	362
18.8	The Just-World Hypothesis . . . . .	364
18.9	Effect Sizes . . . . .	364
18.10	The Contrast Effect . . . . .	366
18.11	How Good—or Bad—are We? . . . . .	368
18.12	Chapter Exercises . . . . .	368

<b>19 Cognitive Dissonance: Psychological Inconsistency</b>	<b>371</b>
19.1 Two Striking Examples . . . . .	372
19.2 Cognitive Dissonance . . . . .	373
19.2.1 How Dissonance Theory Explains the Experiments . . . . .	374
19.3 Insufficient-Justification and Induced-Compliance . . . . .	375
19.3.1 Induced-Compliance and Counter-Attitudinal Behavior . . . . .	375
19.3.2 Prohibition . . . . .	376
19.4 Effort Justification and Dissonance . . . . .	377
19.5 Post-Decisional Dissonance . . . . .	377
19.6 Belief Disconfirmation and Dissonance . . . . .	378
19.6.1 When Prophecy Fails . . . . .	379
19.7 Dissonance Reduction and Bad Reasoning . . . . .	379
19.8 Chapter Exercises . . . . .	380
<b>VIII Evaluating Hypotheses and Assessing Risks</b>	<b>385</b>
<b>20 Causation, Prediction, Testing, and Explaining</b>	<b>389</b>
20.1 Science . . . . .	390
20.1.1 Getting Data . . . . .	391
20.1.2 Testing Hypotheses and Predicting the Future . . . . .	391
20.1.3 Tracking Down Causes . . . . .	393
20.1.4 Mill's Methods . . . . .	396
20.2 Experiments . . . . .	400
20.2.1 Controlled Experiments . . . . .	400
20.2.2 Were the Results due to Chance? . . . . .	402
20.3 Giving Explanations . . . . .	403
20.4 The Everyday Person as Intuitive Scientist . . . . .	405
20.4.1 Gathering Data . . . . .	406
20.4.2 Testing and Predicting . . . . .	407
20.4.3 Tracking Down Causes . . . . .	409
20.4.4 Giving Explanations . . . . .	409
20.5 Intuitive vs. Statistical Prediction . . . . .	409
20.6 Pseudoscience . . . . .	414
20.7 Chapter Exercises . . . . .	417
20.8 Appendix: Scientific Notation and Exponential Growth . . . . .	419
20.8.1 Scientific Notation . . . . .	419
20.8.2 Exponential Growth . . . . .	422

---

<b>21 Risk</b>	<b>427</b>
21.1 Life is Full of Risks . . . . .	428
21.2 Describing Risks . . . . .	429
21.2.1 Risk Ratios . . . . .	429
21.2.2 Exercises . . . . .	431
21.2.3 Finding Information about Risks . . . . .	433
21.3 Health Risks . . . . .	433
21.3.1 The Big Three . . . . .	434
21.4 Crime Risks . . . . .	436
21.4.1 Exercises . . . . .	437
21.5 Other Risks . . . . .	437
21.5.1 Sex . . . . .	437
21.5.2 Love and Marriage . . . . .	437
21.5.3 Jobs and Businesses . . . . .	437
21.6 Cognitive Biases and the Misperception of Risk . . . . .	437
21.6.1 Tradeoffs . . . . .	441
21.7 Psychological Influences on Risk Assessment . . . . .	441
21.7.1 Individuals Differences . . . . .	441
21.7.2 Groups . . . . .	442
21.8 Chapter Exercises . . . . .	442
<b>IX The Social Dimension</b>	<b>445</b>
<b>22 Social Influences on Thinking</b>	<b>449</b>
22.1 The Social World . . . . .	450
22.2 Persuasion: Rational Argument vs. Manipulation . . . . .	450
22.3 Social Influences on Cognition . . . . .	452
22.4 Socialization . . . . .	452
22.5 The Mere Presence of Others . . . . .	453
22.6 Professional Persuaders . . . . .	454
22.6.1 Professional Persuaders: Tricks of the Trade . . . . .	454
22.7 Conformity . . . . .	457
22.7.1 The Autokinetic Effect . . . . .	458
22.7.2 Ash's Conformity Studies . . . . .	458
22.8 Obedience . . . . .	461
22.8.1 The Milgram Experiments . . . . .	461
22.8.2 Changing Behavior vs. Changing Beliefs . . . . .	464
22.9 What could Explain Such Behavior? . . . . .	464
22.9.1 Obedience Training . . . . .	464

22.10	Responsibility . . . . .	465
22.10.1	I Was Just Following Orders . . . . .	465
22.10.2	Asleep at the Wheel . . . . .	465
22.11	Safeguards . . . . .	465
22.11.1	The Open Society and the Importance of Dissent . . . . .	466
22.12	Chapter Exercises . . . . .	467
<b>23</b>	<b>The Power of the Situation</b>	<b>471</b>
23.1	Case Studies . . . . .	471
23.2	The Fundamental Attribution Error . . . . .	477
23.2.1	Explaining why People do what they Do . . . . .	477
23.3	Actor-Observer Differences . . . . .	482
23.4	Special Cases . . . . .	482
23.5	Chapter Exercises . . . . .	484
<b>24</b>	<b>Reasoning in Groups</b>	<b>485</b>
24.1	Group Reasoning . . . . .	486
24.2	Social Loafing . . . . .	486
24.3	Group Dynamics and Setting the Agenda . . . . .	487
24.3.1	Heuristics and Biases in Groups . . . . .	487
24.3.2	Out-group Homogeneity Bias . . . . .	487
24.4	Group Polarization . . . . .	488
24.5	Group Accuracy . . . . .	489
24.6	Groupthink . . . . .	489
24.7	Successful Groups . . . . .	489
24.7.1	Groups in the Classroom . . . . .	489
24.8	Safeguards . . . . .	491
24.9	Chapter Exercises . . . . .	491
<b>25</b>	<b>Stereotypes and Prejudices</b>	<b>493</b>
25.1	Consequences of Prejudices and Stereotypes . . . . .	494
25.1.1	Stereotypes, Prejudice, and Critical Reasoning . . . . .	495
25.2	Prejudice . . . . .	495
25.3	Stereotypes . . . . .	495
25.3.1	Schemas . . . . .	495
25.3.2	Stereotypes as Schemas . . . . .	496
25.4	Discrimination . . . . .	497
25.4.1	Overt vs. Modern Racism . . . . .	498
25.4.2	In-groups vs. Out-groups . . . . .	498
25.5	Features of Stereotypes . . . . .	500

---

25.5.1 Homogeneity . . . . .	500
25.5.2 Resistance to Change . . . . .	501
25.6 Cognitive Mechanisms . . . . .	502
25.6.1 Levels of Generality . . . . .	503
25.7 Flawed Reasoning and Prejudice . . . . .	504
25.8 Responses . . . . .	515
25.8.1 Remedies and Reasoning . . . . .	515
25.9 Chapter Exercises . . . . .	517
<b>26 Social Dilemmas</b> . . . . .	<b>521</b>
26.1 Prisoner's Dilemmas . . . . .	522
26.2 Real- Life Prisoner's Dilemmas . . . . .	524
26.2.1 Public Goods and Free Riders . . . . .	525
26.2.2 The Assurance Problem . . . . .	527
26.3 How is Cooperation Possible? . . . . .	530
26.3.1 Coercion . . . . .	531
26.3.2 Positive By-Products of Cooperation . . . . .	532
26.3.3 Prudence and the Prospect of Future Interaction . . . . .	533
26.3.4 Loyalty . . . . .	534
26.3.5 Moral Principles and Individual Ideals . . . . .	534
26.4 Fundamental Values: Clashes and Tradeoffs . . . . .	535
26.5 Chapter Exercises . . . . .	538
<b>X Representation and Recognition</b> . . . . .	<b>541</b>
<b>27 Diagrammatic Reasoning: Using Pictures to Think</b> . . . . .	<b>545</b>
27.1 Using Vision to Think . . . . .	545
27.2 Picturing Logical Structure . . . . .	547
27.3 Picturing Probabilistic and Statistical Structure . . . . .	553
27.4 Chapter Exercises . . . . .	570
27.5 Appendix: Inverse Probabilities and Bayes' Rule . . . . .	571
<b>28 Recognizing where Cognitive Tools Apply</b> . . . . .	<b>583</b>
28.1 Good Thinking . . . . .	584
28.2 Feedback: Learning from Experience . . . . .	588
28.3 Recognizing the Relevance of a Cognitive Tool . . . . .	589
28.3.1 Recognizing when these Tools are Relevant . . . . .	590
28.3.2 Cues That Signal when a Tool Applies . . . . .	592
28.4 Consider Alternatives . . . . .	593

---

28.4.1 Is there “Invisible Data”? . . . . .	599
28.5 Acquiring Cognitive Skills . . . . .	601
28.6 Chapter Exercises . . . . .	604
<b>Appendix: Pretest</b>	<b>607</b>
<b>Index of Key Concepts</b>	<b>615</b>

# List of Figures

1.1	Going Beyond the Information Given . . . . .	9
1.2	You do your thing, I'll do mine. . . . .	16
4.1	Beyond the Information Given . . . . .	69
4.2	Necker Cube . . . . .	70
4.3	Other Reversing Figures . . . . .	71
4.4	Faces or Vase? . . . . .	71
4.5	An Ambiguous Woman . . . . .	72
4.6	An Ambiguous Animal . . . . .	72
4.7	What's that in the Middle? . . . . .	74
4.8	Two Group Portraits . . . . .	74
4.9	Müller-Lyer Arrow Illusion . . . . .	82
6.1	Viewing a Page in a Browser . . . . .	109
6.2	Searching with AltaVista . . . . .	111
7.1	What's on the Phone? . . . . .	126
7.2	Which Penny is Right? . . . . .	127
7.3	Stages in Memory . . . . .	128
7.4	Classification and Memory . . . . .	130
8.1	Where do False Memories Come From? . . . . .	155
8.2	The Line-up . . . . .	156
13.1	Makeup of a Standard Deck of Cards . . . . .	233
13.2	Outcomes of Rolling Dice . . . . .	233
13.3	A Sentence and its Negation Split One Unit of Probability . . . . .	235
13.4	Negations . . . . .	236
13.5	Disjunctions with Incompatible Disjuncts . . . . .	238
13.6	Tree Representation of the Probability of a Conjunction . . . . .	241

13.7	Table Representation of the Probability of a Conjunction . . . . .	242
13.8	Disjunctions with Compatible (“Overlapping”) Disjuncts . . . . .	243
14.1	Thinning the Relevant Outcomes . . . . .	250
14.2	Conditionalization Trims out a New Unit . . . . .	252
14.3	Tree Diagram of Probability Problem Analysis . . . . .	256
15.1	Inference from Sample to Population . . . . .	279
15.2	Thinking about Correlations . . . . .	289
15.3	Correlation between Smoking and Heart Attacks . . . . .	290
15.4	A Stronger Positive Correlation . . . . .	291
15.5	Independence between Smoking and Heart Attacks . . . . .	291
15.6	Common Causes . . . . .	292
16.1	Feminist Bank Tellers . . . . .	313
18.1	Circles in Context . . . . .	366
18.2	Spin . . . . .	370
20.1	Concomitant Variation . . . . .	398
20.2	Basic Experimental Design . . . . .	401
20.3	Powers of Ten . . . . .	420
20.4	Doubling the Money on Squares . . . . .	423
22.1	Which Line Matches A? . . . . .	459
22.2	Socialization and Conformity . . . . .	460
22.3	Social Influences: The Bad and the Good . . . . .	470
26.1	Prisoner’s Dilemma . . . . .	523
26.2	The Arms Race . . . . .	525
26.3	Taking a Free Ride . . . . .	527
26.4	Playing Chicken . . . . .	540
27.1	Picture of the Structure of an Argument . . . . .	549
27.2	Picture of Necessary and Sufficient Conditions . . . . .	549
27.3	Cards in the Selection Task . . . . .	551
27.4	Testing the Hypotheses in the Selection Task . . . . .	551
27.5	Picturing Validity . . . . .	552
27.6	Probabilities of Disjunctions . . . . .	553
27.7	Probability of Guilt Given Test Results . . . . .	555
27.8	Pie Chart of Lie Detector Test . . . . .	556
27.9	Correlation Diagrams . . . . .	558



---

27.10	Tree for Three Variables . . . . .	559
27.11	Tree Diagram of Wilbur's Lie Detector Test . . . . .	560
27.12	Tree of Possible Outcomes in Monte Hall Problem . . . . .	561
27.13	Two Aces Given at Least One . . . . .	565
27.14	Two Aces Given the Ace of Spades . . . . .	565
27.15	Simple Decision Tree . . . . .	566
27.16	Numbering the Regions . . . . .	567
27.17	Areas Corresponding to $A$ and $B$ . . . . .	567
27.18	Regions Corresponding to Conjunction and Disjunction . . . . .	568
27.19	Diagrams of Negated Compound Sentences . . . . .	569
27.20	Four Valid Argument Patterns . . . . .	569
27.21	Rectangular Diagrams for Disjunctions . . . . .	569



# List of Tables

20.1 Exponential Growth – Powers of Two . . . . .	424
21.1 Causes of Death in America in 1997 . . . . .	430
21.2 Victimization Rates for 1996 . . . . .	436
22.1 Labels on the Shock Generator . . . . .	462
22.2 Results of Milgram Experiments . . . . .	463
23.1 Results of the Good-Samaritan Study . . . . .	472
27.1 Using a Picture to Solve the Monte Hall Problem . . . . .	562
27.2 Simplified Diagram of Monte Hall Problem . . . . .	562
27.3 Diagram for Thinking about Two Aces Problem . . . . .	564



# Introduction

Teaching critical reasoning is difficult. So is learning to reason more carefully and accurately. The greatest challenge is teaching (and learning) skills in such a way that students can spontaneously apply them outside the classroom once the class is over (teaching people to apply them in the classroom can be hard enough, but clearly isn't a worthwhile goal in itself).

I have learned a good deal about these matters from the students who took courses using earlier drafts of this book and from colleagues who taught it. But one key theme of this book is the importance of actually checking to see what the answers to complicated empirical questions are, rather than blithely assuming we know, and that applies to teaching critical reasoning as much as to anything else. As noted in the final chapter, we know more about this now than we did twenty five years ago, but there is still much to learn.

One lesson is clear, though. Reasoning is a *skill*, and there is strong evidence that (like any skill) it can only be acquired with practice. It is also important that students work to apply the concepts and principles in a wide range of situations, including situations that matter to them.

Different routes through the book are possible. One of my colleagues covers virtually all of this book in a single semester. Most of us omit some chapters, however, and the book is designed to accommodate somewhat different courses. A more traditional course would spend a good deal of time on parts two and four (arguments and fallacies), whereas a less traditional course might omit fallacies altogether and focus more on cognitive biases or social aspects of reasoning.

It is also possible to go into probability in more or less detail, although I am convinced that some familiarity with basic probabilistic and statistical concepts is extremely useful for much of the reasoning we commonly do. One can teach this without worrying about calculating a lot of probabilities; indeed, it is important for students to see how the basic concepts apply in cases where precise numbers are unavailable, i.e., in almost all cases they will encounter outside the classroom. Still, having to do some calculations will deepen their grasp of the basic concepts.

Parts of Chapter 20 are under construction but everything else is here. This

book comes in several formats and with several accessories ('crm' stands for critical reasoning manual).

1. crmscreen.pdf. This is a pdf file of the complete manual that is tailored for viewing on a computer. All links in the tables of contents, index, cross references, and so on are live. It is also possible to search the entire manual using the search facilities of the adobe acrobat reader.
2. crmprint.pdf. This is a pdf file of the complete manual tailored for printing (you may need a recent version of adobe acrobat to get a good copy). It will look much better when printed than crmscreen.pdf would.

Both pdf versions require the adobe acrobat reader (preferably version 5.0 or later). You can get it free at:

<http://www.adobe.com/acrobat/readstep.html>

3. crucpostscript.ps. This is a postscript version of the complete manual.
4. This is a pdf copy of the pretest (the pretest is also included as an appendix at the end of the manual). The format is multiple choice. More open-ended questions might yield better information (though it can be difficult to assimilate), but the present format can be graded easily on a scantron, and you can tabulate the results for your class. Then when you discuss a given fallacy or bias (like the conjunction fallacy), you can note how many of the people in the classroom (rather than in some experiment at another university) selected this or that answer.
5. PowerPoint and html versions of slides for a number of chapters are also available.

## **Part I**

# **Basic Concepts of Critical Reasoning**





## **Part I. Basic Concepts of Critical Reasoning**

Most of this module is devoted to a survey of sixteen concepts that will surface repeatedly throughout the course. This will give you some idea what critical reasoning is and what the course will involve. We then turn to issues involving relativism, dogmatism, and the importance of free inquiry.



# Chapter 1

## Basic Concepts of Critical Reasoning

**Overview:** In this chapter we will briefly survey several concepts that will surface repeatedly throughout the book. This will give you some idea what critical reasoning is and what this course will involve.

### Contents

---

<b>1.1 Basic Concepts</b> . . . . .	<b>5</b>
<b>1.2 A Role for Reason</b> . . . . .	<b>13</b>
<b>1.3 Improving Reasoning</b> . . . . .	<b>17</b>
<b>1.4 Chapter Exercises</b> . . . . .	<b>18</b>

---

### 1.1 Basic Concepts

In this section we will briefly survey several concepts that will surface repeatedly throughout the course. This will give you some idea what critical reasoning is and what this course will involve. The aim here is just to provide some basic orientation, so don't worry about details now.

1. Responsibility
2. Reasons
3. Empirical questions

4. Inference and argument
5. Relevance
6. Going beyond the information given
7. The importance of the situation or context
8. Explanation and understanding
9. Prediction
10. Testing
11. Feedback
12. Emotions and needs
13. Quick fixes
14. Persuasion
15. Biases
16. Fallacies
17. Safeguards

We will consider each of these notions briefly (you will find it useful to come back to this list from time to time as you work through later chapters).

In today's rapidly changing world much of what you learn in college will become outdated rather quickly. Many of your grandparents, and perhaps even your parents, had just one or two jobs during their adult life. But the swift pace of globalization and technological innovation make it likely that you will have a succession of jobs, perhaps in quite different fields, once you graduate. Hence it is important for you to *learn how to learn*, and a key part of this is learning how to think critically and carefully about new things.

### **Intellectual Responsibility**

Adults are responsible for the things they do, and this includes thinking clearly and carefully about things that matter. This is hard work and no one succeeds at it completely, but it is part of the price for being in charge of your life.

In addition to thinking for ourselves, it is important to think well. This means basing our reasoning on how things *are*, rather than how we *wish* they were. It means being open to the possibility that we are mistaken, not allowing blind emotion to cloud our thought, and putting in that extra bit of energy to try to get to the bottom of things.

This doesn't mean that we should constantly be questioning everything. Life is too short and busy for that. But in many cases successful action requires planning and thought. It is also desirable to reflect on our most basic beliefs from time to time, and the college years are an ideal time for this. In the end you may wind up

with exactly the same views that you began with. But if you have thought about them carefully, they will be *your own* views, rather than someone else's.

### Reasons

Good reasoning is said to be *cogent*. Cogent reasoning is based on **reasons**. It is based on evidence, rather than on wishful thinking or rash appeals to emotion. When we evaluate a claim our first question should be: What are the *reasons* for thinking it true? If someone tries to convince you to vote for them or that abortion is sometimes the best choice or that God doesn't exist, you should ask: *Why*; what *reasons* are there for thinking that this claim is true?

### Empirical Questions

**Empirical Questions** are questions about what the facts are. They are not matters of opinion, and they are not best answered by guessing. They can only be answered by *checking to see* what the facts are. In the sciences this may involve complex field studies or experiments, but in everyday life the process is often much easier, just one of checking to see. As we will see in various places in the following chapter, answers that seem plausible to us often turn out to be wrong.

### Inference and Argument

When we arrive at a new belief on the basis of reasons, we are said to draw an **inference**. For example, if I learn that 80% of the people in a carefully-conducted poll are going to vote for the Republican candidate for Congress, I might infer (or conclude) that the Republican will win. The results of the poll provide a *reason* to draw this conclusion. If I learn that three of OU's starting five are out with the flu, I may infer (or conclude) that OU will lose to Missouri. My knowledge about the two teams, including the information about the ill players, gives me a *reason* to draw this conclusion.

Such reasoning adds up to an **argument**. Our reasons are the *premises* of the argument, and the new belief is the *conclusion*. For example my inference about the election involves the following argument:

Premise: 80% of the people surveyed plan to vote Republican.  
Conclusion: The Republican candidate for Congress will win.

In a good argument, the premises justify or support the conclusion; they provide good evidence for it.

An argument is a group of sentences; one conclusion and one or more premises. An inference is something we *do* when we draw a conclusion from premises. We will study arguments in detail in the next chapter.

### Relevance

If an argument is to be any good, its premises must be *relevant* to its conclusion. Relevance involves a *relationship* between statements. So a premise can be relevant to one claim while being irrelevant to other claims. It is *irrelevant* if it simply doesn't bear on the truth or falsity of the conclusion, if it's independent of it, if it doesn't affect it one way or the other.

The premise that witnesses claim to have seen Timothy McVeigh rent the Rider truck used in the bombing of the Murrah Federal Building *is* relevant to the conclusion that he is guilty. By contrast, the fact that over a hundred and sixty people were killed in the bombing is *not* relevant to the claim that he's guilty (though once he was convicted it may have been relevant to questions about the appropriate penalty).

One of the major causes of bad reasoning is the use of arguments whose premises are irrelevant to their conclusions. It is very easy to make mistakes about the relevance of one claim to another. This is especially problematic when the premises "look relevant" even though a more careful examination shows that they aren't. Later we will also see that in some cases the acquisition of information of marginal relevance can lead us to dismiss information that is highly relevant to the problem at hand.

### Going Beyond the Information Given

Often our inferences involve leaps from information we are confident about to a conclusion that is less certain. When a pollster conducts a survey to see how the next presidential election is likely to turn out, he asks a few thousand people how they will vote. He then uses this information (about the people in his sample) as a *premise*, and he draws a *conclusion* about what all the voters will do. He has a body of information, what the voters polled say they will do, and he moves beyond this to a conclusion about what voters in general will do.

Our inferences frequently take us beyond the information we already have. Figure 1.1 on the facing page provides a visual representation of this. For example, we often use premises about how things were in the past to draw conclusions about the future. Your doctor relies on her experience when diagnosing your current ailment, and she prescribes a treatment based on what worked best in the past. An experienced wildcatter knows a lot about what geological formations are the best

places to drill for oil and picks a new spot that, he concludes, is a good bet for a well.




Figure 1.1: Going Beyond the Information Given

We also go beyond the information on hand in our personal lives. In the past, people we know have behaved in certain ways, and we frequently conclude that they will behave similarly in the future. Sally has always kept her word, so if I confide in her she probably won't tell anyone; Hank, on the other hand, is a different story. Again, in the past Wilbur had bad experiences going out with people he met in bars, so he concludes that this isn't a good way for him to meet people and looks around for alternatives.

When we draw a conclusion that goes beyond the information we have, there is always a risk that we'll be wrong. But we will see that if we use certain strategies, we can increase the likelihood that we will be right. In some cases we can use numbers to measure just how likely this will be. This means that in the chapters on probability you will have to manipulate just a few fractions, though nothing more than what you did in Algebra 1 in high school.

Inferences that go beyond the information that we have are pervasive; indeed, in Part III we will see that even perception and memory often go beyond the information given in much the way that many inferences do. When our inference carries us beyond information we are sure about, we always run the risk of being wrong, but we will discover some strategies that will reduce this risk.

### The Importance of Context

Reasoning, inference, and decision making never occur in a vacuum. We will see over and over again that the **context** or **situation** in which we think about things can strongly influence the *ways* in which we think about them. Indeed, it even affects how we perceive and remember things. Furthermore, our reasoning is sometimes

faulty because we underestimate the importance of context. We will see that this is especially true when we are trying to understand the behavior of other people.

### Explanation and Understanding

**Explanation reflex:** we have a strong need to understand and make sense of the world around us

We are constantly trying to make sense of things. We need to *explain* and *understand* the world around us. Almost every time we ask *why* something happened or *how* something works we are seeking an explanation.

Learning about things and understanding how they work is often rewarding in and of itself, and it is vital if we are to deal successfully with the world around us. If we understand how things work, we will be able to make more accurate predictions about their behavior, and this will make it easier for us to influence how things will turn out. If you understand how an automobile engine works, you will be in a much better position to fix it the next time it breaks down. If you understand basic principles of nutrition, you will be in a better position to lose weight and keep it off.

We are constantly seeking explanations in our daily lives. My computer worked yesterday; everything seems the same today, so what explains the fact that it won't boot up now? We are particularly concerned to understand the behavior of other people. *Why* did Clinton lie to the Grand Jury about his affair with Monica Lewinsky—and why did he have the affair in the first place? *Why* did Timothy McVeigh bomb the Alfred P. Murrah Federal Building? *Why* did the people in Heaven's Gate so happily commit suicide? Such questions also arise closer to home. "*Why* did Sally give Wilbur that look when he said they should go out again; what did she mean by it?" In fact, we often have occasion to wonder why *we* do some of the things that we do; "*Why* in the world did I ever say such an idiotic thing as *that*?"

We are always looking for reasons and regularities and patterns in the phenomena around us. Much reasoning involves attempts to explain things, and sometimes leads us to see patterns that are not really there or to accept overly simplified explanations just to have the feeling that we understand what is going on.

For example, some things really do happen by coincidence. But it can be tempting to seek an explanation for them, for example, to adopt some superstition to account for things that just happened by chance. Again, people who like conspiracy theories want simple, pat explanations for why things are going badly for them. When we later learn what really happened in such cases (like the Watergate coverup), we often find less subtle and intricate conspiracy than we imagined, and more bungling and accident. But a conspiracy would offer such a nice simple explanation of things. So one goal in later chapters will be to devise good explanations while avoiding bad ones.



## Prediction

We use reasoning to **predict** what will happen. If I tighten the bolts, the garden gate will probably last for another year. If I tell Sam what I really think about the ghastly color of his new car, he'll go ballistic. When we make predictions, we use the information that seems *relevant* to us (e.g., information about Sam and his short temper) and draw an *inference* about what *will* happen. We will see that there are common patterns of errors that can arise in this process.

## Testing

Our beliefs are much more likely to be true if they are *based on evidence*. It isn't enough for a scientist to just propose a new theory. The theory has to be **tested**, and it needs to survive stringent tests. We typically test a theory by using it to make a prediction, and we then see if the prediction comes true. If it does, that provides some (though by no means conclusive) support for the theory; if it does not, the theory is in trouble. For example, the germ theory of disease was only accepted once it had been used to make a variety of successful predictions, e.g., once vaccines were shown to be effective.

Science works as well as it does because it is responsive to evidence in this way. And our views in daily life will also be more likely to be true if we test them. We will see, however, that most of us aren't very good at this.

## Feedback

Testing our ideas is one way of getting feedback. Without **feedback** telling us how accurate our reasoning has been we won't be able to learn from our mistakes. Feedback is often painful; we learn that we didn't do as well as we had thought or hoped—maybe we didn't do very well at all. But reasoning, like so much else, involves trial and error, and unless you know what the errors are, you won't do any better the next time around. So if we want to improve our ability to reason and make judgments we must seek feedback.

We often overlook the importance of feedback. For example, people who conduct job interviews may have a good deal of experience. Even so they typically receive limited feedback on their hiring abilities. Why? They do feedback about the quality of the people that they hire, but they rarely get feedback on the quality of the people they reject.

### Emotions and Needs

**Emotions** are a central part of our lives, and they often play a quite legitimate role in our thinking. Intense emotions, however, can lead to poor reasoning. If we are extremely frightened or extremely angry, we aren't likely to think very clearly. Less obviously, emotions often provide an *incentive* to think badly. For example, the desire to avoid unpleasant facts about ourselves or the world can lead to wishful thinking and to various self-serving biases in our thought.

We cannot be effective thinkers if we won't face obvious facts or if we seriously distort them. Good thinking involves reasoning, not rationalization; it is based on what we have *good reasons* to think is true, not on what we would *like* to be true. Throughout this course we will see how desires and emotions and moods can impair clear thinking, and we will discuss ways to minimize their effect.

### Quick Fixes

We encounter many difficult problems in today's world: the rise in terrorism, racism and racial tensions, the growing sense that jobs are not secure, the increasing pollution of the environment all present huge challenges. On a more personal level, the desire to save a marriage or to quit smoking or to make more money also present challenges. In such cases genuine solutions are likely to require a great deal of time or effort or money (or, often enough, all three), and in some cases it isn't even clear where to begin.

The solutions to these problems like this would require us to do things we don't want to do. Most of us don't want to spend a lot of our own money to solve the problems of the world or adopt a new lifestyle even though it's healthier. So it is not surprising that people who promise us a quicker and easier solution—a quick fix—will always find an audience. A **quick fix** is something that is offered as a fast and easy solution to a complex problem. The human tendency to wishful thinking is one reason why claims by those who offer a quick fix are often accepted, even when there is little *evidence* in their favor. We will find that hopes for a quick fix are responsible for a good deal of careless reasoning.

### Persuasion

We often try to **persuade** others to accept our view or position. People in the “persuasion professions” (like advertising, politics, charity work) do this for a living, but we all do it some of the time. You might want to convince someone to go out with you, or to marry you, or to give you a divorce.

There are many different (and often subtle) techniques for persuading people. Some involve offering them reasons; others rely on manipulations. We have noted

that people prefer having bad reasons to no reasons, and so manipulation often works best if it is disguised to look like an argument. We want to think that reasons and arguments can be given to support our views, even if those arguments aren't very good. As a result, one very effective way of persuading people is to appeal to their emotions (e.g., their self-interest or their fears) but to dress the appeal up as an argument that doesn't appear to appeal to their feelings.

We will encounter various techniques for persuasion throughout the course. Some involve good arguments. Some (called fallacies) masquerade as good arguments (when they really aren't). We will also examine various non-rational ways of persuading people. If we are aware of these, we will be less likely to be taken in by them.

### **Biases**

**Biases** are systematic tendencies to reason badly. We will encounter a number of biases in the following chapters. All of us are vulnerable to biases, but understanding how they work and seeing how pervasive they are will help us to minimize their influence in our own thinking and to spot their results in the thinking of others.

### **Fallacies**

Bad reasoning is said to be **fallacious**. If our reasoning is biased we are likely to commit fallacies. In Part Four of the course we will study several common fallacies.

### **Safeguards**

Throughout the course we will learn various safeguards for counteracting common biases in thought and avoiding fallacies.

## **1.2 A Role for Reason**

Many important issues seem very difficult to settle. We may wonder whether *any* observations or research or arguments could show that physician-assisted suicide or abortion or capital punishment or a flat income tax or the use of marijuana are right (or wrong, as the case may be). There are three ways to avoid wrestling with such difficult issues.

1. We can simply refuse to think about such questions at all.
2. We can embrace some view that furnishes quick and simple answers to these questions.

3. We can decide that such questions cannot be answered on the grounds that beliefs about things like morality and religion are simply subjective.

We will consider these in turn.

### **I Don't want to Think about it**

If we adopt the first option, simply refusing to think about difficult topics, we drift through life like robots. This isn't a good way to live; one way to see this is to ask yourself whether you want to raise your children so that they turn out this way. Moreover, difficult decisions sometimes have to be made, and it is a good thing for us to have a voice in those decisions that affect us. Finally, there are cases where we simply can't avoid making a hard decision, cases where tuning out and doing nothing itself has terrible consequences.

### **Dogmatism: The True Believer**

The dogmatist is a *true believer* in some theory or doctrine. The key feature of the true believer is not what he believes, but *how* he believes it. The true believer is not open minded; he wouldn't let anything count as evidence against his beliefs. The true believer's views provide a set of principles and categories, and he interprets things in terms of them.

It is possible to be dogmatic about all sorts of things. Countless people have been dogmatic about their political views. Many Marxists, Nazis, and others were so certain of their views that they were willing to murder millions of people to translate their theory into practice. One can also be dogmatic about religious views.

True believers tend to see things as all black or all white, and so they often think that most questions have simple answers. They are often uncompromising, and sometimes feel that those who disagree with them are not just wrong, but evil, an enemy that must be conquered. Sometimes the resistance to the enemy is peaceful, but history shows that it can also be very bloody.

### **Faith**

Some beliefs that are not based on reason and evidence are based on *faith*. Faith surely has its place, but *in the absence of reason* it can be used to justify *anything*. The faith of those who truly believe that they belong to the master race can be just as strong as the faith of anyone else. Members of a Satanic cult that practices human sacrifice may be just as certain that they are right as anyone ever is. Taken just by itself, faith can justify anything.

### Relativism: Who is to Say?

In the face of such difficulties, some people adopt the view that values are *subjective*. It's "all relative;" "who's to say" what's right and wrong? According to this view, issues of morality are like issues of etiquette. Many people agree that we shouldn't eat peas with a knife and, similarly, many agree that we should help those in need. But this is really just an opinion shared by people in *our* culture or society. There are no objective facts about such things, and other societies might, with equal legitimacy, adopt quite different views.

Relativism can seem appealing because it offers an easy answer to the difficult questions about right and wrong: we don't have to wrestle with them, because they are subjective, simply matters of opinion or taste. It can also seem an attractively tolerant view. Live and let live: we have our views, other groups have theirs, and since there is no fact of the matter about what is right we should just leave each other alone.

It is possible to extend this relativistic stance to things besides values. Indeed, an extreme relativist might claim that *everything* is relative, that there are no objective facts at all. But this more extreme version of relativism is incoherent. If everything is relative, then the very claim "everything is relative," is relative too. The claim "everything is relative" can be true for some people and false for others, and there is no fact of the matter about which group of people is correct. The claim undercuts itself.

Relativism about values doesn't collapse as easily as the more extreme relativism, but it sounds much better in theory than it does in practice. It is easy to say things like "Well, that's a question about values, and those are just subjective," but very few of us would accept the following implications of this view.

1. Since there is no objective answer to questions about right and wrong, it's just a matter of opinion that I should not grade the finals for this course by tossing them down the steps and basing the grade for each exam on the distance it travels before landing.
2. Since everything is relative, it doesn't really matter what my children grow up believing.
3. Since values are subjective, who are we to say that flying airliners into the twin towers of the World Trade Center is wrong?
4. Since values are relative, varying from one culture to another, who is to say that the Nazis were wrong to kill millions of Jews? It might be wrong-for-us, but it was right-for-them.



Figure 1.2: You do your thing, I'll do mine.

### Tolerance and Open-mindedness

Relativism may sound like a nice, tolerant view, but it really isn't. The claim that we should tolerate others is itself a claim about values, and it cannot be defended by the claim that there really are no (objective) values. Worse, the consistent relativist must grant that intolerant societies (like Nazi Germany) are no more wrong about things than any other society. In the end, according to the relativist, it's just a matter of taste or opinion whether tolerance is a good thing or not.

The claim that there are objective truths does not mean that we have cornered the market on them. In some cases the truth may be unknown, and in some cases other cultures may be much nearer the mark than we are. But relativism wouldn't allow this. If others can't be wrong, then they also can't be right, and so there really isn't any sense in which we can learn from them.

Many years ago I saw a poster like figure 1.2 in a "head shop"; it makes the point about relativism better than any abstract discussion ever could.

### Fallibility: Commitment with an Open Mind

Most people are neither full-fledged true believers or full-fledged relativists. There are various intermediate positions, but one that fits especially well with a commitment to free, independent, critical reasoning is called *fallibilism*.

The fallibilist believes that virtually all of our views are fallible. Almost any of them could turn out to be false. But this does *not* mean that all our beliefs are equally well supported or equally good. A fallibilist acts on the *best* reasons and evidence she can get, while remaining open minded and willing to change her views if new evidence or arguments require doing so. This often means living with uncertainty, but that is just the human condition.

A free society that is open to dissent makes critical reasoning much easier. In a society where open discussion is allowed, different viewpoints can be aired. Without free expression the scope of our thoughts will be limited; we will be exposed to fewer novel ideas, and our sense of the range of possibilities will be constricted. Since no one has cornered the market on truth, we should beware of those who would set themselves up as censors to decide what the rest of us can say and hear.

### 1.3 Improving Reasoning

Critical reasoning is a *skill*, and like all skills it requires active involvement. As you read this book you will learn how to use various *intellectual* or *cognitive tools* (e.g., some logical and probability rules, various rules of thumb, diagrams) that will help you reason better. But as with all skills, you can only learn by practice. If you just passively read the chapters or absorb lectures you will not learn anything worth learning. You can only master these tools by using them in a variety of contexts (including outside the classroom). There are no foolproof rules that will always lead to good reasoning, but there are three things that will improve your thinking.

1. Be aware of the most common ways in which reasoning can go wrong; this will help you guard against them in your own thinking and spot them in the thinking of others.
2. Use the rules of thumb (discussed in the rest of the course); this will make it easier for you to reason well.
3. Try to apply the things you learn in the course outside the classroom.

The third step is the hardest, but it is vital. Many of our actions result from habit, and our habits of relying on past views and acting without really thinking are a chief cause of defective reasoning. Even once we master the material in this course, it is easy to lapse back into automatic pilot once we leave the classroom. It is difficult, in the midst of a busy life, to simply stop and think. But that is what we have to do if we are going to think more carefully about the things that matter to us most.<sup>1</sup>

---

<sup>1</sup>To avoid cluttering the text with footnotes references will be given at the end of each chapter (in some cases the references have yet to be added). The first study involving giving reasons when there weren't any is by R. Nisbett & T. D. Wilson, "Telling More than we can Know: Verbal Reports on Mental Processes," *Psychological Review*, 84 (1977); 231–259. The Xerox machine study was conducted by E. Langer, A. Blank, & B. Chanowitz, "The Mindlessness of Ostensibly Thoughtful Action: The Role of "Placebic" Information in Interpersonal Interaction," *Journal of Personality and Social Psychology* 3 (1978); 365–342. Further references to be supplied.

## 1.4 Chapter Exercises

1. Suppose that you had been adopted at birth by a family very different from your own and that you had been raised in a very different subculture or even in a very different country.
  - a. Do you think that your political views would have been different? If not, why not? If so, how? Give two concrete examples of beliefs you find important in your own life that you might not have had.
  - b. Do you think that your views about right and wrong would have been different? If not, why not? If so, how? Give two concrete examples of beliefs you find important in your own life that you might not have had.
  - c. Do you think that your religious beliefs would have been different? If not, why not? If so, how? Give two concrete examples of beliefs you find important in your own life that you might not have had.
2. Give an example of a view you once thought was obviously true but which you now think might be false. What led you to change your mind?
3. Give an example of a view you once thought was obviously false but which you now think might be true. What led you to change your mind?



## **Part II**

# **Reasons and Arguments**



## Part II. Reasons and Arguments

An *argument* is a claim that is backed by reasons. In the Chapter 2 we study the nature of arguments, learn ways to identify them in their natural habitat, and meet two key notions, deductive validity and inductive strength.

In Chapter 3 we learn about one very important kind of sentence, the conditional. Conditionals are *iffy*; they tell us that if one thing is true, then something else will be true as well. We will also learn about necessary and sufficient conditions and study four very important kinds of conditional arguments.



# Chapter 2

## Arguments

**Overview:** An argument is a claim that is backed by reasons. In this chapter we study the nature of arguments, learn ways to identify them in their natural habitat, and encounter two key concepts, deductive validity and inductive strength.

### Contents

---

<b>2.1</b>	<b>Arguments</b> . . . . .	<b>24</b>
2.1.1	Inferences and Arguments . . . . .	24
<b>2.2</b>	<b>Uses of Arguments</b> . . . . .	<b>25</b>
2.2.1	Reasoning . . . . .	25
2.2.2	Persuasion . . . . .	25
2.2.3	Evaluation . . . . .	25
<b>2.3</b>	<b>Identifying Arguments in their Natural Habitat</b> . . . . .	<b>26</b>
2.3.1	Indicator Words . . . . .	26
<b>2.4</b>	<b>Putting Arguments into Standard Form</b> . . . . .	<b>27</b>
2.4.1	Arguments vs. Conditionals . . . . .	27
<b>2.5</b>	<b>Deductive Validity</b> . . . . .	<b>30</b>
2.5.1	Definition of Deductive Validity . . . . .	30
2.5.2	Further Features of Deductive Validity . . . . .	33
2.5.3	Soundness . . . . .	33
<b>2.6</b>	<b>Method of Counterexample</b> . . . . .	<b>34</b>
<b>2.7</b>	<b>Inductive Strength</b> . . . . .	<b>39</b>
<b>2.8</b>	<b>Evaluating Arguments</b> . . . . .	<b>41</b>

---

2.9 Chapter Exercises . . . . . 42

---

## 2.1 Arguments

### 2.1.1 Inferences and Arguments

We draw an *inference* when we make a judgment based on some evidence or assumptions or reasons. You learn that 87% of the people in a carefully-conducted poll are going to vote for the Republican candidate for Congress, and you infer (or conclude) that the Republican will win. The results of the poll provide a *reason* to draw this conclusion. Inference is an activity or process, something we *do* when we draw a conclusion from assumptions or premises. By contrast, arguments (as we will use the term) are not processes but groups of sentences. Still, we can often study and evaluate inferences by looking at the argument patterns they involve.

You learn that 87% of the people plan to vote Republican and conclude the Republican candidate will win. This inference or reasoning follows the contours of an argument. Our reasons are the *premises* of the argument, and the new belief is the *conclusion*. Thus our inference about the election involves the following argument:

Premise: 80% of the people surveyed plan to vote Republican.  
 Conclusion: The Republican candidate for Congress will win.

In a good argument, the premises justify or support the conclusion; they provide good evidence for it.

An **Argument** consists of two things:

1. a group of one or more sentences – the **premises**
2. one further sentence – the **conclusion**

Any time someone gives reasons to support a claim, they are giving an argument. Their argument *makes a claim*; this claim is its conclusion. The premises are intended to provide reasons (justification, support, evidence) for the conclusion. Some arguments are good and some aren't. In a good argument, the premises really do provide a good reason to think that the conclusion is true. Sometimes we also call premises reasons, or assumptions.

The sentences that make up an argument must all be ones that can be either true or false. So an argument consists of declarative (or indicative) sentences (rather than commands or questions). In everyday life we often think of an argument as a

**Argument:**

1. One or more Premises
2. One Conclusion

dispute or disagreement in which people shout at each other. But for purposes of critical reasoning an argument is just a group of declarative sentences: one of them is the **conclusion**; the rest are **premises**.

## 2.2 Uses of Arguments

Arguments can be used for various purposes. You will understand the various functions of arguments better once you have analyzed a number of examples, but we will mention three of the main functions of arguments here.

### 2.2.1 Reasoning

We often use arguments when we are engaged in *problem solving* or *deliberation*. We reason about what would happen if we did certain things or if certain events occurred. In such situations we are not trying to justify a particular claim. We are instead interested in what would follow if certain things *were* true.

This is sometimes called *what if* reasoning. What, for example, would follow if you spend \$3500 on a car this fall? Well, let's assume that you do. Let's do a little calculation. That would only leave you \$450. Not good. So what would follow if you spend \$2800? You try out different possibilities and reason about their consequences.

### 2.2.2 Persuasion

We often use arguments in an attempt to convince or persuade someone else that something is true. In such cases, we are trying to get someone to accept the conclusion of our argument by giving them reasons (premises) to believe it. Our audience may be huge (as with the viewers of a presidential debate or the readers of an editorial in a large newspaper) or small (maybe just one other person—or even just oneself). But in each case, the aim is to make a definite claim and to *justify* it by giving reasons that support it.

### 2.2.3 Evaluation

Often we need to *evaluate* other people's arguments, and sometimes to *refute* them. Is their argument any good? And if it isn't, it may be important to be able to say just where it goes wrong. In such cases we need to reason critically in order to show that the premises of the argument do not adequately support its conclusion.

Some of the inferences we draw are good and others are not. Some arguments are strong and others are weak. A chief goal of a book like this one is to help you

separate the good from the bad. This will be a goal throughout most of the book, but first we need to see how to identify arguments when we come across them.

## 2.3 Identifying Arguments in their Natural Habitat

The first step in deciding whether something that you read or hear is an argument is to determine whether it has a **conclusion** and (if so) what that conclusion is. Once you identify the conclusion, you can usually figure out the premises. So you should *always begin by looking for the conclusion*.

### 2.3.1 Indicator Words

Some words are good indicators of a conclusion; the typical job of these words is to say *Here comes the conclusion*. To discover what these words are, consider a very simple argument:

**Premise 1:** All humans are mortal.

**Premise 2:** Socrates is a human.

**Conclusion:** \_\_\_\_\_ Socrates is mortal.

The last sentence of this argument is its conclusion. To discover some conclusion indicators, just ask yourself what words could sensibly be placed in the blank before the last sentence (think about this and write down several before proceeding). Some natural words and phrases to put in the blank are ‘therefore’, ‘hence’, ‘and so’, ‘thus’, ‘consequently’, and ‘it follows that’.

Sometimes the conclusion of an argument comes at the beginning, rather than at the end. In this case any indicator words will come after the conclusion. To see how this works, we’ll just rearrange our argument a bit.

**Conclusion:** Socrates is mortal, \_\_\_\_\_

**Premise 1:** All humans are mortal.

**Premise 2:** Socrates is human.

What words can sensibly go in this blank (write some down before proceeding)? Some natural choices are ‘because’, ‘for’, and ‘since’. These words usually say *Here comes a premise*.

Conclusion Indicators

**Conclusion Indicators** ‘therefore’, ‘thus’, ‘so’, ‘hence’, ‘consequently’, ‘accordingly’, ‘entails that’, ‘implies that’, ‘we may conclude that’, ‘this establishes that’, ‘this gives us reason to suppose that’, ‘in short’, ...



**Premise Indicators** ‘because’, ‘for’, ‘since’, ‘after all’, ‘inasmuch as’, ‘in view of the fact that’, ‘in virtue of’, ‘here are the reasons’, ...

These lists are not exhaustive, but they include the key indicators. In most cases our rules of thumb for identifying premises and conclusions, based on the occurrence of these indicator words, work well. But there are exceptions, so our guidelines are *never* a substitute for *thinking about the example yourself*.

If arguments came neatly packaged and labeled like the two we’ve just seen, things would be easy. But in real life arguments often do not contain any indicator words at all. When this happens ask yourself: what is the other person trying to get me to believe? Once you figure this out, you’ll have the conclusion, and then it should be relatively easy to locate the premises.

In real life things can be even more unclear. Sometimes a premise of an argument isn’t explicitly stated; other times the conclusion is missing. There are often good reasons why parts of the argument aren’t stated. Sometimes they are obvious enough in the context that they “go without saying.”

One or more *unstated premises* are common when the premises include information that is widely known, obvious, or easily figured out in the context. And an *unstated conclusion* often occurs when the conclusion is thought to follow so obviously from the premises that it would be insulting to your intelligence to have to spell it out for you.

## 2.4 Putting Arguments into Standard Form

Although real-life arguments rarely come neatly packaged, it will make our task easier if we adopt a clear format for *repackaging* them. We will say that an argument is in *standard form* if it consists of a list of all the premises followed by the conclusion. In an argument with two premises, it will take the following form:

1. All Oklahomans are Sooners.
  2. Tom is an Oklahoman
- Conclusion: Tom is a Sooner.

The line above the conclusion makes it easy to take in at a glance what the conclusion is. If there are more than a couple of premises it is often handy to number them, though in simpler cases it isn’t necessary.

### 2.4.1 Arguments vs. Conditionals

One type of sentence, conditionals, are easily confused with arguments. Here is an example of a conditional sentence.

- If Monica had thrown away her blue dress, Clinton would have escaped impeachment.

This is *not* an argument. It does not give reasons to support any claim. It doesn't advance any conclusion. It just says that *if* one thing is the case, *then* something else is too. But this is merely hypothetical. You could assert this, but then go on to add, correctly, that she didn't throw away the dress and that he did not escape impeachment. Statements that make hypothetical claims like this are called *conditionals*. Such statements are easily confused with arguments, but they do not contain premises or conclusions. So they aren't arguments.

Compare the above sentence with

- Monica did throw away her blue dress. Therefore, Clinton did escape impeachment.

This *is* an argument (as it stands it's not a very good argument, but bad arguments are arguments too). It makes two *claims*, and neither is hypothetical at all. It says (falsely) that Monica did throw away the infamous dress. It also says (again falsely) that Clinton did escape impeachment. The word 'therefore' is a conclusion indicator, and here it signals that the claim that Clinton escaped impeachment is the conclusion of the argument.

When you are trying to determine whether a passage contains an argument, ask yourself "What is the other person's *point*; what is she trying to get me to believe?" "What if anything is being supported and what is just merely being asserted?"

In real life the conclusion will often come at the very end of an argument. But it can also occur at the beginning. It may even come in between various premises. It may also be the case that some of the material is extraneous, padding or filler that really isn't part of the argument at all.

### Exercises on Identification of Arguments

Which of the following passages expresses an argument (remember: an argument may not be a very good argument). For each one that does

- Enclose its conclusion in parentheses
- Underline each premise
- Draw a wavy line under any extraneous material.
- If the passage doesn't contain an argument, write "NA" beside it (if you can also say *why* it isn't an argument, that's even better).
- In some cases you may have to supply a missing premise or conclusion.

1. Clinton asked Vernon Jordan to help get Lewinsky a job. That amounts to trying to buy her testimony. So he is guilty of obstruction of justice.
2. Although Clinton did some terrible things, they don't rise to the level of impeachment.
3. As long as members of al-Qaeda are on the loose, we will be in danger.
4. The Zamori tribe will eventually die out, because they initiate their young by putting them to death at the age of four (George Carlin).
5. Kevin grew up in Oklahoma, so he knows what Oklahomans need most from a meteorologist (Channel 4).
6. The race is not always to the swift or the victory to the strong, but that's the way to bet (Dayman Runyan).
7. The ability to hold the breath is important in swimming, since it develops confidence and permits the practice of many swimming skills.
8. The House vote to impeach Clinton was pretty much along party lines. So the entire exercise was pretty partisan.
9. Intramural basketball started recently, so the fitness center will be crowded because lots of games will be going on.
10. Blessed are the meek: for they shall inherit the earth (New Testament).
11. Blessed are the meek, but they ain't gonna get rich (J. R. Ewing).
12. The prime numbers can't come to an end. If they did, we could multiply all the prime numbers together and add one, and this would give us a new number. But this new number would be prime, because it would leave a remainder of 1 when divided by any of the prime numbers.
13. The Iowa State Lottery has already contributed over 41 million dollars to Iowa's massive economic development push. The funds have been used all around the state: business incentives, cultural program, agricultural research, trade and export development, education. When you play the lottery, all of Iowa wins.
14. If you are sure, from the betting and the draw, that a player has three of a kind and you have two pairs, even aces over kings, you are a sucker to go for the full house, because the odds against buying that ace or king are prohibitive at 10.8 to 1.

15. Consciously or unconsciously, the reader is dissatisfied with being told only what is not; she wishes to be told what is. Hence, as a rule, it is better to express even a negative claim in positive form.
16. Somebody else will get AIDS, but not me.
17. You can't argue about right and wrong, because it's all a matter of opinion anyway. Who's to say what's right and wrong?
18. Shut the door—if you don't want the dog to get in.
19. If Osama bin Laden planned the attacks, we must track him down. And it's clear that he did. So we've got to find him.

## 2.5 Deductive Validity

### 2.5.1 Definition of Deductive Validity

In this section we will learn about deductive validity. It is easy to define this notion, but it is deceptively abstract and slippery. It takes practice to master it.

**Validity:** *If* all premises true,  
conclusion *must* be true

When the premises of an argument support its conclusion in the strongest possible way, we say that the argument is **deductively valid**. There are several different, but equivalent, ways to define deductive validity:

1. A deductively valid argument is one such that *if* all of its premises are true, its conclusion *must* be true.
2. A deductively valid argument is one such that it is *impossible* for its conclusion to be false when all of its premises are true.

The most common mistake to make about validity is to think that this definition says more than it actually does. It does not say anything whatsoever about the premises (taken in isolation from the conclusion) or about the conclusion (taken in isolation from the premises). Deductive validity involves the *relationship* between the premises and the conclusion. It says that a certain *combination* of the two, all true premises and a false conclusion, is impossible.

True premises in  $\Rightarrow$  true  
conclusion out

Deductively valid arguments are *truth preserving*; if all of them are true, then the conclusion must be true as well. True premises in, true conclusion out. Since there is only one sort of validity, namely deductive validity, we will often speak simply of 'validity'.

### Understanding Deductive Validity

If you will spend a few minutes thinking about the following examples, you will begin to get a feeling for what validity really means. Suppose someone makes the following two claims and that you believe they are true:

1. If it is raining, then the parking lot will be full.
2. It is raining.

Now ask yourself. What can I conclude from 1 and 2? The answer isn't difficult, but it is important to reflect on it.

3. The lot will be full.

Note that you do *not* need to know whether premises 1 and 2 are true in order to see that the claim that *the lot is full* follows from them. It is this notion of *following from* that means that this argument is deductively valid.

Now ask yourself: is there *any* consistent, coherent story that we could imagine in which 1 and 2 were both true, but in which the claim that the lot is full was false? Try it. It can't be done. If you try it and begin to see that it's impossible, you are on your way to understanding deductive validity.

Here's a second example:

1. If Wilbur won the race, he would have called to brag about it.
2. But he hasn't called.

What follows from this? Is there any possible way that sentences 1 and 2 could both be true while at the same time the sentence

3. Wilbur did not win.

is false?

What about:

1. If a set is recursive, then it is recursively enumerable.
2. The set you mention is recursive.

What follows from this? Note that you don't even need to understand the words 'recursive' or 'enumerable', much less know whether these two sentences are true, to see that sentence 3 follows logically from them.

3. The set you mention is recursively enumerable.

By way of contrast, consider the following argument:

1. If Wilbur won the race, he would have called to brag about it.
2. He did call to brag.

If we know that 1 and 2 are both true, can we be sure that

3. Wilbur won the race.

No, we can't be sure. In this case, it *is* possible for the two premises to be true while the conclusion is false. For example, Wilbur might be a compulsive liar who called to brag even though he came in last.

### **Validity: Less is More**

As noted above, it is a very common mistake to think that the definition of deductive validity says *more* than it actually does. It *only* says what has to be the case *if* all of the premises are true.

1. The definition does *not* require that either the premises or the conclusion of a valid argument be true.
2. The definition does *not* say anything about what happens if one or more of the premises is false. In particular, it does *not* say that if any of the premises are false, then the conclusion must be false.
3. The definition does *not* say anything about what happens if the conclusion is true. In particular, it does *not* say that if the conclusion is true, then the premises must be true.

The definition of validity only requires that the premises and conclusion be related in such a way that *if* the premises are (or had been) true, the conclusion is (or would have been) true as well.

The definition of deductive validity is hypothetical; *if* all of its premises are true, *then* the conclusion must be true as well. The *if* here is a big one. It's like the *if* on the postcard you get that announces: "You will receive ten million dollars from the Publishers' Clearing House—if you hold the winning ticket." This does not mean that you do have the winning entry. And similarly it does not follow from the definition of deductive validity that all of the premises of each deductively valid argument are true.

There can be deductively valid arguments with:

1. False premises and a false conclusion

2. False premises and a true conclusion
3. All true premises and a true conclusion

The *only* combination that *cannot* occur in a deductively valid argument is all true premises and a false conclusion. This can never happen, because by definition a deductively valid argument is one whose form makes it *impossible* for all of its premises to be true *and* its conclusion false.

Invalid arguments can have any of these three combinations plus the combination of all true premises and a false conclusion (which is the one combination that valid arguments cannot have).

### 2.5.2 Further Features of Deductive Validity

1. Deductive validity *does not come in degrees*. It is all or none.
2. In a deductively valid argument the *conclusion contains no new information*; there is no information in the conclusion that was not already contained in the premises.

We won't worry about the second feature now, but it will become important when we turn to inductively strong arguments.

### 2.5.3 Soundness

An argument is *sound* just in case

1. It is *deductively valid*, and
2. It has *all true premises*.

Soundness = Validity + all true premises

Once you master the concept of validity (which is tricky), soundness will be easy. The conclusion of a sound argument must be true. Why?

#### A Note on Terminology

- Only arguments can be valid or invalid; sentences or statements cannot.
- On the other hand, only statements or sentences can be true or false (have a truth value); arguments can be neither.

## 2.6 Method of Counterexample

We can use the *method of counterexample* to show that an argument is *invalid*. The method consists of telling a *consistent* story in which all the premises of the argument are true, but the conclusion is false.

**Counterexample:** a consistent scenario in which all premises are true but conclusion is false

The idea behind the method is this. It is *impossible* for a deductively valid argument to have all true premises while having a false conclusion. A counterexample is a possible scenario in which the premises are true and the conclusion is false. So it shows that it *is* possible for the argument to have all true premises and a false conclusion. And this proves that it is *deductively invalid*.

Consider the argument:

1. If Jones drove the getaway car, he's guilty of the robbery.
  2. Jones did not drive the getaway car.
- 
- So Jones is not guilty of the robbery.

Counterexample: Suppose that Jones did not drive the getaway car, but he *was* one of the other robbers. In this case both premises would be true and the conclusion false. This means the argument is invalid.

Now here is one for you to try:

1. If Tom overslept, he will have been late to work.
  2. Tom was late to work \_\_\_\_\_.
- So Tom did oversleep.

Can you construct a counterexample to show this argument deductively invalid?

### Exercises on Validity

Answers to selected exercises are given on page [37](#).

1. What conclusions are obvious consequences of the following sets of premises (the first one is worked for you, as an example)?

- A.     If Tom is from Texas, then Will is from Alabama.  
        2. Tom is from Texas.  
        \_\_\_\_\_
- So Will is from Alabama.

- B.     1. If Tom is from Texas, then Will is from Florida.  
        2. Will is not from Florida.  
        \_\_\_\_\_



- C.     1. Either Sara is from Texas or she is from Florida.  
       2. Sara is not from Florida.

2. Which of the following arguments are valid? In many cases you won't know whether the premises are true or not, but in those cases where you do, say whether the argument is sound.

- A.     1. All Republicans hate the poor.  
       2. Newt Gingrich is a Republican.  
       So, Newt Gingrich hates the poor.

- B.     1. All Democrats cheat on their spouses.  
       2. All men are Democrats.  
       Therefore, all men cheat on their spouses.

- C.     1. If my alarm breaks, I'm late to work.  
       2. I was late to work.  
       Therefore, my alarm broke.

- D.     1. If my alarm breaks, I'm late to work.  
       2. I was not late to work.  
       Therefore, my alarm did not break.

- E.     1. Many Fords run for years without any problems  
       2. My car is a Ford.  
       Therefore, my car will run for years without any problems.

- F.     1. Gore and Lieberman can't both win the Democratic nomination.  
       2. Lieberman will be nominated.  
       Therefore, Gore won't get the nomination.

- G. 1. OU and OSU can't both win the Big 12 outright.  
 2. OU will not win the Big 12 outright.  
 Therefore, OSU will win the Big 12 outright.
- H. 1. If we don't have free will, we can't be blamed for our actions.  
 2. If we can't be blamed for our actions, we shouldn't be punished.  
 Therefore, If we don't have free will, we shouldn't be punished.
- I. 1. If Sam committed first degree murder, then he intended to kill Tom.  
 2. And he did intend to kill him (he admitted it in his testimony.)  
 So Sam is guilty of murder in the first degree.
3. Use the method of counterexample to show that the following arguments could each have all true premises while having a false conclusion. If you succeed, this will prove that they are invalid.
- A. Some politicians are honest.  
Will is a politician.  
 So Will is honest.
- B. If Jane Fonda is the U.S. President, then she is famous.  
Jane Fonda is famous.  
 So Jane Fonda is President.
- C. Some Democrats are not honest people.  
 So some honest people are not Democrats.
- D. Whenever Bill is home his car is in the garage.  
Bill's car is in the garage.  
 So Bill must be home.
4. Construct a valid argument with at least one false premise.
5. Construct a valid argument with a false conclusion.
6. Is it possible to construct a valid argument with all true premises and a false conclusion? If not, why not?
7. Describe the method of counterexample.
8. When it's properly employed what does the method of counterexample show?

**Answers to Selected Exercises**

1. What conclusion are obvious consequences of the following sets of premises?

B.     1. If Tom is from Texas, then Will is from Florida.  
       2. Will is not from Florida.  
       So, Tom is not from Texas

C.     1. Either Sara is from Texas or she is from Florida.  
       2. Sara is not from Florida.  
       So Sara is from Texas.

2. You were asked which of these arguments were valid. Explanations are given below the relevant argument.

A.     1. All Republicans hate the poor.  
       2. Newt Gingrich is a Republican.  
       So, Newt Gingrich hates the poor.

Argument A. is valid, but it is unsound because the first premise is false.

B.     1. All Democrats cheat on their spouses.  
       2. All men are Democrats.  
       So all men cheat on their spouses.

The argument is valid, but it is unsound because both premises are false.

C.     1. If my alarm breaks, I'm late to work.  
       2. I was late to work.  
       Therefore, my alarm broke.

Invalid; it would be possible for both premises to be true and the conclusion false. This would happen if the two premises were both true but I was late for some other reason, e.g., my car didn't start or it broke down on the way to work.

D.     1. If my alarm breaks, I'm late to work.  
       2. I was not late to work.  
       Therefore, my alarm did not break.

Valid; this one is harder. Drawing a picture will help. But don't worry a lot about it yet; We'll study arguments like this in Chapter 3.

- E.     1. Many Fords run for years without any problems  
        2. My car is a Ford.  
        Therefore, my car will run for years without any problems.

Invalid; the premises could be true but the conclusion false. This would be the case if my car were one of the exceptions, one of the few Fords that were lemons.

- G.     1. OU and OSU can't both win the Big 12 outright.  
        2. OU will not win the Big 12 outright.  
        Therefore, OSU will win the Big 12 outright.

Invalid; the premises could both be true and the conclusion could still be false. This would be the case if any of the ten teams not from Oklahoma won. It would be the case, for example, if Nebraska won outright.

H. Valid.

I. Invalid.

3. Use the method of counterexample to show that the following two arguments could each have all true premises while having a false conclusion. If you succeed, this will prove that they are invalid.

- A.     Some politicians are honest.  
        Will is a politician.  
        So Will is honest.

Here is one counterexample (there are many others). Imagine that Will is a politician who accepts bribes from his constituents. In this case the premises are true, but the conclusion is false. This provides a counterexample that shows this argument is invalid.

5. Construct a valid argument with a false conclusion.

Both the arguments about Republicans and the argument about Democrats are valid (4). And each of them has at least one false premise, so both of them also have a false conclusion (5).

6. Is it possible to construct a valid argument with all true premises and a false conclusion? If not, why not?

This is impossible. The *definition of validity* does not allow this possibility.

8. What does the method of counterexample show (when properly employed)?

The method yields a possible scenario in which all of the premises of an argument are true and its conclusion is false. This provides a counterexample to the claim that the argument is deductively valid, thereby showing that the argument is *invalid*.

## 2.7 Inductive Strength

We will study inductive strength in detail later on, so we will just briefly consider it here.

**Inductive strength:** if premises all true, a high probability conclusion is true

An argument is *inductively strong* just in case:

1. It is not deductively valid, and
2. If all of its premises are true, then there is a *high probability* that its conclusion will be true as well.

The second item is the important one. The only point of the first item is to insure that no argument is both deductively valid *and* inductively strong (this will make things easier for us in later chapters).

There are two important ways in which Inductive strength differs from deductive validity:

1. Unlike deductive validity, inductive strength *comes in degrees*.
2. In a deductively valid argument, the conclusion does not contain any information that was not already present in the premises. By contrast, in an inductively strong argument, the *conclusion contains new information*

A deductively valid argument with all true premises *must* have a true conclusion. By contrast, an inductively strong argument with true premises provides good, but not conclusive grounds, for its conclusion.

Since we have defined things so that inductively strong arguments are not deductively valid, we can think of arguments as arranged along a continuum of descending strength:

1. Deductively valid

2. Deductively invalid
- (a) Inductively strong
  - ⋮
  - (b) Inductively weak
  - ⋮
  - (c) Worthless

### General and Particular

It is sometimes said that deductively valid arguments proceed from the general to the specific, whereas inductively strong arguments proceed from the specific to the general. This is not a good way to think about the two sorts of arguments, and we have defined things so that notions of generality and specificity are completely irrelevant to the two notions. Here is a deductively valid argument that goes from more specific premises to a more general conclusion:

- 1. 3 is a prime number.
  - 2. 5 is a prime number.
- 
- Therefore all odd numbers between 2 and 6 are prime.

And here is an inductively strong argument that goes from a more general premise to a more specific conclusion.

- 1. All the crows observed thus far have been black.
- Hence, the next crow to be observed will be black.

### Deductively Valid and Inductively Strong Reasons

Sometimes it is more natural to speak of reasons rather than arguments. A group of sentences provides deductively valid reasons for a conclusion just in case it is impossible for all of them to be true and the conclusion false. Valid reasons have this feature because there is no information in the conclusion that was not already contained in the reasons themselves.

**Inductively strong reasons:** A group of sentences provides inductively strong reasons for a conclusion just in case it is unlikely for all of them to be true and the conclusion false. If a group of inductively strong reasons for a conclusion are true, then there is a good chance that the conclusion will be true as well, but there is also some possibility that it will be false. Inductively strong reasons are not always truth preserving. There is an *inductive leap* from the reasons to the conclusion. Inductive support comes in varying degrees; the stronger the inductive reasons, the less risky the inductive leap.

## 2.8 Evaluating Arguments

We do much of our reasoning almost automatically, so it is easy to overlook how frequently we engage in it. Any time someone gives reasons to support a claim, they are giving an argument. Much of this course is devoted to the evaluation of arguments, and we will find three key issues that surface over and over again.

Once you have identified an argument, you must ask three questions.

1. Are the premises true (or at least plausible)?
2. Has any relevant information been omitted from the premises?
3. Do the premises support the conclusion?

**1. Are the premises plausible?** If the issue is whether you should believe the conclusion, then the first question to ask is whether the premises are plausible. In this context, nothing can salvage an argument if one or more of its premises are false. If the premises of an argument—even just one of them—are false, we have no reason to accept its conclusion.

Sometimes we can't be certain whether the premises of an argument are true, and we have to settle for plausibility instead. But the more plausible the premises, the better. If you are simply doing *what if* reasoning, the plausibility of the premises may not matter; you are just asking what would be true *if* the premises were true, and in this context it doesn't matter whether they actually are true.

**2. Has relevant information been omitted?** When it comes to reasoning, ignorance is not bliss; what you don't know *can* hurt you. An argument may have all true premises and yet omit information that is relevant to our evaluation of it. Suppose Wilbur tells me that Jack would be a good person to buy a used car from because Jack knows a lot about cars and doesn't use high-pressure techniques. These premises may be true, but if Wilbur fails to mention that Jack has done time for fraud, I'm in trouble if I accept the conclusion of his argument.

We can't usually get the whole truth and nothing but the truth; examining all of the evidence that might conceivably be relevant would be an endless task. But we should never neglect evidence that we know about or evidence that seems like it might bear on the issue in a major way.

In many cases we need a good deal of background knowledge to answer the first two questions. If the argument is about football, we need to know something about football; if it is about cooking, we need to know something about cooking. Logic can't supply this information, but we will discuss various things, e.g., evaluation of sources, that can help us answer the first two questions when they arise in real life.

We will see that this question doesn't arise when we are evaluating a deductively valid argument.

**3. How strongly, if at all, do the premises support the conclusion?** This is another way of asking whether the argument is deductively valid or inductively strong and, if it's the latter, just how strong it is.

In order to master various key concepts, we will sometimes focus on one of these questions without worrying about the others. But when we put things all together at the end, when you are evaluating reasoning in the real world, all three questions are important.

## 2.9 Chapter Exercises

The first two passages contain arguments: (a) say whether the argument is valid; (b) if it isn't valid, say whether it is inductively strong; (c) if it is inductively weak, say why. Give a careful analysis of the overall strength of the argument.

1. Over 96% of all Fords sold in 1996 had to go back into the shop in the two year period from 1997 to 1999. Wilbur's bought his Ford in 1996. So it was supposed to go back into the shop.
2. Most of the Fords sold in 1995 had to go back into the shop in the two year period from 1997 to 1999. Wilbur's bought his Ford in 1995. So it was supposed to go back into the shop.
3. If an argument is inductively strong does it have to be the case that if the conclusion is true, then the premises are very likely to be true?

**Answers to Selected Exercises** 3. No, this gets it backwards. If an argument is inductively strong then if the premises are true the conclusion is likely to be true, but this doesn't mean that if the conclusion is true, then premises are very likely to be true. Example:

1. Pat is wearing a frilly pink dress.  
Hence, Pat is a female.

This argument is inductively strong; if the premise is true, the conclusion probably is as well. But this no longer holds if we swap premise for conclusion:

1. Pat is a female.  
Hence, Pat is wearing a frilly pink dress.



## Chapter 3

# Conditionals and Conditional Arguments

**Overview:** In this chapter we will learn about one very important kind of sentence, conditionals, and the central role they play in reasoning. Conditionals are *iffy*; they tell us that if one thing is true, then something else will be true as well. We will also learn about necessary and sufficient conditions and study four very important kinds of conditional arguments.

### Contents

---

<b>3.1</b>	<b>Conditionals and their Parts</b>	<b>43</b>
3.1.1	Alternative Ways to State Conditionals	45
<b>3.2</b>	<b>Necessary and Sufficient Conditions</b>	<b>48</b>
<b>3.3</b>	<b>Conditional Arguments</b>	<b>53</b>
3.3.1	Conditional Arguments that Affirm	53
3.3.2	Conditional Arguments that Deny	54
<b>3.4</b>	<b>Chapter Exercises</b>	<b>55</b>

---

### 3.1 Conditionals and their Parts

A **conditional** is a sentence that says something will be true *provided that* something else is. All of the sentences are conditionals:

1. If it is raining, then the lot is full.
2. If Clinton lied to the Grand Jury, then he should be tried for perjury.
3. If it doesn't fit, you must acquit.
4. If Congressman Gary Condit had nothing to hide, he wouldn't be so sneaky.
5. If he builds it, they will come.
6. If you bomb the final, you'll fail the course.
7. If a number is divisible by 2, then it isn't a prime number.

A conditional is a *compound* sentence that consists of two shorter sentences. When the sentence has an *if-then* format,

1. the sentence between the *if* and the *then* is the **antecedent**
2. the sentence after the *then* is the **consequent**.

The antecedent is the part that comes before (think about Poker, where you ante up before the hand begins). The antecedent of sentence 1. is

- It is raining

and the consequent is

- The lot is full.

The words *if* and *then* are *not* part of either the antecedent or the consequent. They are just connecting words that glue the two simpler sentences together to form the conditional.

What conditionals do *not* say

A conditional is *iffy*. It does *not* claim that its antecedent *is* true or that its consequent *is* true. It is *hypothetical*: *if* the antecedent is true, *then* the consequent will be true too. As we noted above, it is true that *if* you win the Publishers Clearing House Sweepstakes, *then* you will be rich. But unfortunately this does not mean that you will win the sweepstakes or that you will be rich.

Our first example of a conditional:

1. If it is raining, then the lot is full.

tells us what will happen *if* it is raining (the lot will be full). The biggest trouble people have with conditionals is thinking that they say *more* than they do.

1. Sentence 1. does *not* say anything about what happens if the lot is full. (In particular, it does not say that if the lot is full, it will be raining.)
2. Sentence 1. does *not* say anything about what happens if it is not raining. (In particular, it does not say that if it's not raining then the lot isn't full.) Our definition of validity is naturally stated as a conditional: *if* all of the premises are true, *then* the conclusion must be true as well. And one of the main problems people have with it is thinking that it says more than it does.

### 3.1.1 Alternative Ways to State Conditionals

There are various ways to state conditionals, and some of them require thought. Ask yourself: could we *rephrase* a sentence as an *if-then* claim without changing its meaning? If we can, then the sentence is a conditional.

The following sentences all have the same meaning, so we count all of them as conditionals.

1. If it is raining, then the lot is full
2. If it is raining, the lot is full.
3. When it's raining, the lot is full.
4. The lot is full, if it's raining.
5. It rains, and the lot is full.
6. Should it rain, the lot will be full.
7. The lot is full, providing that it's raining.

#### A Note on Terminology

1. All arguments have premises and a conclusion. But no argument has an antecedent or a consequent.
2. All conditionals have antecedents and consequents. But no conditional has a premise or a conclusion.

#### Exercises on Conditionals

Determine whether or not each of the following sentences is a conditional. If it is, draw a solid line under its antecedent and a broken line under its consequent. Answers to selected exercises are given on page 46.

Example: If Gore wins the next presidential election, then we'll have a democrat for president in the early years of the next century.

1. This sentence is a conditional.
2. Antecedent: Gore wins (or will win) the next Presidential election
3. Consequent: We'll have a democrat for president in the early years the next century.

Note: the words 'if' and 'then' are not part of either the antecedent or the consequent.

1. If we run out of gas out here in the desert, we're as good as dead.

2. If you have the time, we have the beer.
3. If Betty goes to the movies, Tom will stay home and watch the kids.
4. Betty will stay home and watch the kids, if Tom goes to the movies.
5. Tom will stay home and watch the kids, provided that Betty calls him on time.
6. When it rains, it pours.
7. Should OU win the rest of their games, they'll win the Big 12 Conference.
8. Give peace a chance
9. Give me a place to stand, and a tall frosty Bud Light.
10. Give me a place to stand, and I'll move the world.
11. Tom went to the movies and Betty watched the kids.
12. Wilbur will fix your car only if you pay him what you owe.
13. Wilbur will fix your car if only you pay him what you owe.
14. If OU wins their next game, then if they win the one after that, they'll have had a good season.

#### **Answers to Selected Exercises on Conditionals**

1. If we run out of gas out here in the desert, we're as good as dead.
  - (a) This is a conditional.
  - (b) Antecedent: We do run out of gas in the dessert.
  - (c) Consequent: We're as good as dead.
2. If you have the time, we have the beer.
  - (a) This is a conditional.
  - (b) Antecedent: You have the time
  - (c) Consequent: We have the beer
3. If Betty goes to the movies, Tom will stay home and watch the kids.
4. Betty stayed home and watched the kids, if Tom went to the movies.

- 
- (a) This is a conditional
  - (b) Antecedent: Tom went to the movies
  - (c) Consequent: Betty stayed home and watched the kids
5. Tom will stay home and watch the kids, provided that Betty calls him on time.
  6. When it rains, it pours.
    - (a) This is a conditional.
    - (b) Antecedent: It rains.
    - (c) Consequent: It pours.
  7. Should OU win the rest of their games, they'll win the Big 12 Conference.
  8. Give peace a chance
    - Not a conditional
  9. Give me a place to stand, and a tall frosty Bud Light.
    - This is not a conditional. In fact, it's not even a declarative sentence.
  10. Give me a place to stand, and I'll move the world.
    - (a) Although this looks a lot like the previous sentence, this *is* a conditional. You have to focus on what it means. The key is to see that it says the same thing as "If you give me a place to stand, then I'll move the world."
    - (b) Antecedent: You give (gave) me a place to stand.
    - (c) Consequent: I'll move the world.
  11. Tom went to the movies and Betty watched the kids.
    - Not a conditional
  12. Wilbur will fix your car only if you pay him what you owe.
    - (a) This is a conditional.
    - (b) Antecedent: Wilbur will fix your car.
    - (c) Consequent: You pay him what you owe.
  13. Wilbur will fix your car if only you pay him what you owe.

- (a) This is a conditional.
  - (b) Antecedent: You pay him what you owe.
  - (c) Consequent: Wilbur will fix your car. [Note the difference between this and the preceding problem]
14. If OU wins their next game, then if they win the one after that, they'll have had a good season.
- (a) This is a conditional.
  - (b) Antecedent: OU will win their next game
  - (c) Consequent If they win the one after that, they'll have a good season.  
Note that the consequent of sentence 12. is itself a conditional.

### 3.2 Necessary and Sufficient Conditions

**Sufficient condition:** enough, a guarantee

One sentence is a **sufficient condition** for a second if the truth of the first would *guarantee* the truth of the second. The truth of the first is *enough*—all you need, *sufficient*—to ensure the truth of the second. Having your head cut off is a sufficient condition for being dead. There are many other ways to die, but decapitation is enough to ensure it.

**Necessary condition:** a condition that must be met

And one sentence is a **necessary condition** for a second if the truth of the second sentence is required—is needed, is *necessary*—for the truth of the first sentence. For example, paying your tuition fees is a necessary condition for graduating. If you are to graduate, you must pay your fees.

In the previous example, paying your fees is necessary, but not sufficient, for graduating. It is necessary, because no one graduates without paying their fees, but it is not sufficient because there are other things you also have to do (like pass the required number of credit hours).

Requirements and prerequisites are necessary conditions

Requirements and prerequisites are usually necessary conditions. They are things that you must do in order to achieve a certain goal, but by themselves they do not guarantee success. For example, studying is a necessary condition for passing this course (you must study to pass it), and practicing is a necessary condition for becoming a good basketball player (you have to practice). But neither, alone, is sufficient.

1. The **antecedent** of a conditional is a **sufficient** condition for the **consequent**.
2. The **consequent** of a conditional is a **necessary** condition for the **antecedent**.

We have treated sentences as necessary and sufficient conditions, but it is also useful to think of properties in this way. For example, the property or characteristic

of being a dog is sufficient for that of being an animal. Having the first property is enough to insure having the second. And the property or characteristic of being an animal is necessary for that of being a dog. Nothing is a dog unless it is an animal.

Only sentences can serve as antecedents or consequents of conditionals, but we can often turn talk about properties into talk about sentences. For example, instead of saying that the property of being a dog is sufficient for the property of being an animal we could say that if something is a dog then it is an animal. This involves subtleties you will learn about if you take a course in logic, but they aren't relevant here. When we do this, do antecedents are still sufficient conditions for their consequents and consequents are still necessary conditions for their antecedents.

The best way to understand these concepts is to work the following exercises on necessary and sufficient conditions.

### Exercises on Necessary and Sufficient Conditions

Answers to selected exercises are given on page 51.

1. "If you get a cholera shot, then you'll be safe in the villages." Here the claim is that getting a cholera shot is a \_\_\_\_\_ condition for being safe in the villages.
2. "You must get a cholera shot, if you're going to get a visa to India." Here the claim is that getting a cholera shot is a \_\_\_\_\_ condition for getting a visa to India.
3. "Lots of people with rich friends do not succeed in national politics. But no one succeeds in national politics without rich friends." Here the claim is that having rich friends is a \_\_\_\_\_ condition (but not a \_\_\_\_\_ condition) for succeeding in national politics.
4. "Show me a good loser, and I'll show you a loser." Here the claim is that being a good loser is a \_\_\_\_\_ condition for being a loser.
5. "Someone is a bachelor if, and only if, they are an unmarried male." Here the claim is that being an unmarried male is a \_\_\_\_\_ condition for being a bachelor.
6. "OU will go to a Bowl game only if they have a winning season." Here the claim is that having a winning season is a \_\_\_\_\_ condition for going to a Bowl game.

7. "OU won't go to a Bowl game unless they have a winning season." Here the claim is that having a winning season is a \_\_\_\_\_ condition for going to a Bowl game.
8. "...liberal democracy has arisen only in nations that are market oriented, not in all of them, but only in them" [Charles E. Lindblom, *Politics and Markets*]. Here the claim is that a society's being market oriented is a \_\_\_\_\_ (but not a \_\_\_\_\_) condition for liberal democracy.
9. "If wishes were horses, then beggars would ride." Here the claim is that wishes being horses is a \_\_\_\_\_ condition for beggars to ride.
10. "Hard work guarantees success." Here the claim is that working hard is a \_\_\_\_\_ condition for being successful.
11. "The defendant is only guilty of first-degree murder if she planned the crime out beforehand." Here the claim is that planing the crime out beforehand is a \_\_\_\_\_ condition for being guilty of first-degree murder.
12. "I'm not about to run in this heat." Here the claim is that the heat is a \_\_\_\_\_ condition for *not* running.
13. "Show me someone who likes your cooking, and I'll show you someone who needs a tongue transplant." Here the claim is that liking your cooking is a \_\_\_\_\_ condition for needing to get their tongue replaced.
14. "A contract is binding only when there is no fraud." Here the claim is that the absence of fraud is a \_\_\_\_\_ condition for a contract to be binding.
15. "If Hank is a father then he is a male." Here the claim is that being male is a \_\_\_\_\_ condition for being a father.
16. "Nothing ventured, nothing gained." Here the claim is that taking a chance is a \_\_\_\_\_ condition for gaining something.
17. "You must pass the final to pass this course." Here the claim is that passing the final is \_\_\_\_\_ condition for passing the course.
18. "You will pass this course only if you pass the final." Here the claim is that passing the final is a \_\_\_\_\_ condition for passing the course.
19. "You will pass this course if only you pass the final." Here the claim is that passing the final is a \_\_\_\_\_ condition for passing the course. (This sentence is tricky; compare this sentence to the one right above it.)



20. “You will not pass this course unless you pass the final.” Here the claim is that passing the final is a \_\_\_\_\_ condition for passing the course.
21. “John will marry Sue only if she agrees to have three children.” Here the claim is that agreeing to have three children is a \_\_\_\_\_ condition for his agreeing to marry her.
22. “You won’t be happy, if you buy it at Sturdleys.” Here the claim is that buying it at Sturdley’s is a \_\_\_\_\_ for being unhappy.
23. “A person is a brother just in case he is a male sibling.” Here the claim is that being a male sibling is a \_\_\_\_\_ condition for being a brother.

### Answers to Selected Exercises on Necessary and Sufficient Conditions

1. “If you get a cholera shot, then you’ll be safe in the villages.” Here the claim is that getting a cholera shot is a sufficient condition for being safe in the villages.

The claim here is that if you get a shot, you’ll be safe. So the sentence says that getting a shot is *enough*, it’s sufficient for being safe.

2. “You must get a cholera shot, if you’re going to get a visa to India.” Here the claim is that getting a cholera shot is a necessary condition for getting a visa to India.

Requirements and prerequisites are usually necessary conditions. Here the claim is that you must get the shot if you are to get the visa. But there are other necessary conditions as well, e.g., not being a convicted felon. So getting the shot is not sufficient.

3. “Lots of people with rich friends do not succeed in national politics. But no one succeeds in national politics without rich friends.” Here the claim is that having rich friends is a necessary (but *not* a sufficient condition) for succeeding in national politics.

This sentence says that a requirement, a necessary condition, for success in national politics is having rich friends. But since lots of people with rich friends do not succeed, being rich is not sufficient.

4. “Show me a good loser, and I’ll show you a loser.” Here the claim is that being a good loser is a sufficient condition for being a loser.

Hint: it will help here to rephrase the claim in a more explicitly conditional form, as “If you show me a good loser, I’ll show you a loser.”

5. “Someone is a bachelor if, and only if, they are an unmarried male.” Here the claim is that being an unmarried male is a *both* a necessary and sufficient condition for being a bachelor.
6. “OU will go to a Bowl game only if they have a winning season.” Here the claim is that having a winning season is a necessary condition for going to a Bowl game.
7. “OU won’t go to a Bowl game unless they have a winning season.” Here the claim is that having a winning season is a necessary condition for going to a Bowl game.
8. “. . . liberal democracy has arisen only in nations that are market oriented, not in all of them, but only in them” [Charles E. Lindblom, *Politics and Markets*]. Here the claim is that a society’s being market oriented is a necessary (but *not* a *sufficient* condition) for liberal democracy.
9. “If wishes were horses, then beggars would ride.” Here the claim is that wishes being horses is a sufficient condition for beggars to ride.
10. “Hard work guarantees success.” Here the claim is that working hard is a sufficient condition for being successful.
11. “The defendant is only guilty of first-degree murder if she planned the crime out beforehand.” Here the claim is that planing the crime out beforehand is a necessary condition for being guilty of first-degree murder.
12. “I’m not about to run in this heat.” Here the claim is that the heat is a sufficient condition for *not* running.
13. “Show me someone who likes your cooking, and I’ll show you someone who needs a tongue transplant.” Here the claim is that liking your cooking is a sufficient condition for needing to get their tongue replaced.
14. “A contract is binding only when there is no fraud.” Here the claim is that the absence of fraud is a necessary condition for a contract to be binding.

15. “If Hank is a father then he male.” Here the claim is that being male is a necessary condition for being a father.
16. “Nothing ventured, nothing gained.” Here the claim is that taking a chance is a necessary condition for gaining something.
17. “You must pass the final to pass this course.” Here the claim is that passing the final is necessary condition for passing the course.
18. “You will pass this course only if you pass the final.” Here the claim is that passing the final is a necessary condition for passing the course.
19. “You will pass this course if only you pass the final.” Here the claim is that passing the final is a sufficient condition for passing the course. (Compare this sentence to the one right above it.)

### 3.3 Conditional Arguments

#### 3.3.1 Conditional Arguments that Affirm

Arguments that have a conditional as one premise and either the antecedent *or* the consequent of that very conditional as the second premise are called **conditional arguments**. The first type of conditional argument we will study has the antecedent of the conditional as the second premise.

1. If bin Laden planned the attacks, we must track him down.
  2. He did plan the attacks.
- 
- Therefore (3) We must track him down.

The first premise of this argument is a conditional and the second premise says that the *antecedent* of that conditional is *true*. The second premise just repeats—affirms—the antecedent in the first premise. We say that such arguments **affirm the antecedent**. All arguments that affirm the antecedent are *deductively valid*. It is impossible for an argument with this format to have all true premises and a false conclusion. This format is sometimes known by its Latin name *modus ponens*.

By contrast, in the argument

1. If Norman is in Oklahoma, then Norman is south of Kansas.
  2. Norman is south of Kansas.
- 
- Therefore (3) Norman is in Oklahoma.

Affirming the antecedent  
—valid

Affirming the consequent  
–invalid

the first premise is a conditional and the second premise says that the *consequent* of the conditional is *true*. Such arguments **affirm the consequent**. *Each and every* argument that has this format is *deductively invalid*. It is possible for such arguments to have all true premises and a false conclusion. Affirming the consequent is *always* a fallacy.

### 3.3.2 Conditional Arguments that Deny

#### Negations

We have studied one kind of sentence, the conditional. Now we need to introduce another kind, the negation. The **negation** of a sentence is another sentence which says that the first sentence is false. It says the *opposite* of what the first sentence says; it *denies* it. We could express the negation of the sentence

It is raining.

by any of the following sentences:

1. It is not the case that it is raining.
2. It is not true that it is raining.
3. It is not raining.
4. It isn't raining.
5. It ain't raining.
6. Ain't rainin'.

Arguments that have a conditional as one premise and either the negation of that conditional's antecedent *or* the negation of the conditional's consequent are also **conditional arguments**. So there are two conditional arguments that affirm and two more that deny, for a total of four. Here is a conditional argument in which the second premise is the negation of the antecedent of the first premise.

In the argument

1. If Norman is in Oklahoma, then Norman is south of Kansas.
2. Norman is not south of Kansas.

Therefore (3) Norman is not in Oklahoma.

Denying the consequent  
–valid

the first premise is a conditional and the second premise says that the *consequent* of the conditional is *false*. Such arguments **deny the consequent**. *Each and every* argument that has this format is *deductively valid*. This format is sometimes known by its Latin name *modus tollens*.

By contrast, in the argument

1. If Norman is in Oklahoma, then Norman is south of Kansas.

2. Norman is not in Oklahoma.

Therefore (3) Norman is not south of Kansas.

the first premise is a conditional and the second premise says that the *antecedent* of the conditional is *false*. Such arguments **deny the antecedent**. All arguments having this format are *deductively invalid*. Denying the antecedent is *always* a fallacy. Denying the antecedent  
–invalid

Here are two examples:

If he builds it, they will come. But they didn't come. So he didn't build it.

We repackage the argument in standard form like this:

1. If he builds it, they will come.

2. They didn't come.

Therefore (3) He didn't build it.

It is impossible for both of premises of this argument to be true while its conclusion is false, and so is deductively valid. The argument denies the consequent.

If the sawdust is the work of carpenter ants, then we'll need something stronger than Raid to fix the problem. But fortunately it's not the work of carpenter ants, so we won't need anything stronger than Raid.

In standard form:

1. If the sawdust is the work of carpenter ants, then we'll need something stronger than Raid.

2. The sawdust is not the work of carpenter ants.

Therefore (3) We won't need anything stronger than Raid.

This argument commits the fallacy of denying the antecedent. Hence it is invalid. But we should be able to see this without knowing the label: *if* you knew that the two premises were true, you still could not be sure whether the conclusion was true or not. The sawdust might be the work of termites (in which case we'll definitely need something stronger than Raid.)

### 3.4 Chapter Exercises

Answers to selected exercises are given on page 58.

1. Put each of the following conditional arguments into standard form. Then say which form the argument has and whether it is valid or invalid. Remember that we have four types of conditional argument:

- Affirming the antecedent (always valid)
  - Affirming the consequent (always invalid)
  - Denying the antecedent (always invalid)
  - Denying the consequent (always valid)
1. If Stan had gotten the job, he would have called me to brag about it by now. But he hasn't called. So he didn't get it.
  2. If Sara passed the bar exam, she will call me to brag about it by now. And she did pass it.
  3. If Sara passed the bar exam, she would have called me to brag about it by now. But she didn't pass it.
  4. OU played in a Bowl game last year only if they had a winning season. And I hate to tell you, but they didn't have a winning season.
  5. You will do well in Critical Reasoning if you keep up with the assignments, and you do keep up.
  6. My car only dies when the temperature is below freezing. But it's below freezing today, so it will die.
  7. If Smith embezzled the money, then Jones was involved in the crime. But Smith didn't embezzle it. So Jones wasn't involved.
  8. If Wilbur is an uncle, he couldn't have been an only child. And he is an uncle.
  9. If wishes were horses, then beggars would ride. But beggars don't ride. So it looks like wishes aren't horses.
  10. If everyone here at OU were doing well with the current course requirements, there would be no need to change the requirements. But some people are not doing well.
  11. If OU has a winning record in the Big 12, then if all of their players are healthy, they will do well in the tournament. And they have a winning record.
  12. If Tom's prints are on the gun, then he is guilty. So he must be innocent, because those weren't his prints on the weapon.
  13. If morals could be taught simply on the basis that they are necessary to society, there would be no social need for religion. But morality cannot be taught in that way – Patrick Lord Devlin, *The Enforcement of Morals*.

14. Wilbur is guilty of first degree murder only if he intended to kill the victim. But he was in such a rage he couldn't really have intended anything. So he isn't guilty in the first degree.
  15. The jury must vote not guilty if they have a reasonable doubt about the guilt of the defendant. And they can't help but have a reasonable doubt in this case.
  16. Suppose that you have a pack of special cards, each of which has a letter [either a consonant or a vowel] on one side and a number [either even or odd] on the other. If you have some of the cards lying flat on a table, which ones should you turn over in order to determine whether cards with vowels on one side always have odd numbers on the other side (this exercise is harder)?
    - (a) cards with consonants and cards with even numbers on them.
    - (b) cards with vowels and cards with even numbers on them.
    - (c) cards with consonants and cards with odd numbers on them.
    - (d) cards with vowels and cards with odd numbers on them.
    - (e) you need to turn over *all* of the cards, in order to determine whether or not this is so.
    - (f) None of the above.
2. What is the relationship between sufficient conditions and the rule of affirming the antecedent? What is the relationship between necessary conditions and the rule of denying the consequent?
3. Here are a few more exercises on necessary and sufficient conditions.
1. "You will not pass this course unless you pass the final." Here the claim is that passing the final is a \_\_\_\_\_ condition for passing the course.
  2. "John will marry Sue only if she agrees to have three children." Here the claim is that agreeing to have three children is a \_\_\_\_\_ condition for his agreeing to marry her.
  3. Vixens are female foxes. Being a vixen is a \_\_\_\_\_ condition for being a female fox?
  4. "You won't be happy, if you buy it at Sturdleys." Here the claim is that buying it at Sturdley's is a \_\_\_\_\_ for being unhappy.
  5. "A person is a brother just in case he is a male sibling." Here the claim is that being a male sibling is a \_\_\_\_\_ condition for being a brother.

4. If  $A$  is a sufficient condition for  $B$ , then the negation of  $B$  is also a sufficient condition for the negation of  $A$ . For example, being a dog is a sufficient condition for being an animal. And not being an animal is a sufficient condition for not being a dog. Explain why this holds in general, and draw a diagram to illustrate your points.

5. If  $A$  is a necessary condition for  $B$ , then the negation of  $B$  is also a necessary condition for the negation of  $A$ . For example, being an animal is a necessary condition for being a dog. And not being a dog is a necessary condition for not being an animal. Explain why this holds in general, and draw a diagram to illustrate your points.

### Answers to Selected Chapter Exercises

1. The direction here were to put each conditional arguments into standard form, to say which form the argument has, and then to say whether it is valid or invalid.

1. If Stan had gotten the job, he would have called to brag about it by now. But he hasn't called. So he didn't get it.

1. If Stan had gotten the job, he would have called to brag about it by now.

He hasn't called.

3. Therefore, he didn't get it.

The conclusion is that Stan didn't win the race. The argument denies the consequent, so it is valid.

2. If Sara passed the bar exam, she will call me to brag about. And she did pass.

1. If Sara passed the bar exam, she will call me to brag about it.

She did pass it.

3. Therefore, she will call me to brag about it.

The conclusion of this argument isn't included in it. You have to supply it. The conclusion is that she called me to brag. The argument affirms the antecedent, so it is valid

3. If Sara passed the bar exam, she would have called me to brag about it by now. But she didn't pass it.

Invalid. Why?



4. OU played in a Bowl game last year only if they had a winning season. And I hate to tell you, but they didn't have a winning season.

The first sentence here is a premise, but it is a little tricky. It says that *if* OU played in a Bowl, they had a winning season. The next sentence says that they didn't have a winning season, so it denies the consequent of the first sentence. You have to supply the conclusion (what is it?). Since the argument denies the consequent, it is valid.

5. You will do well in Critical Reasoning if you keep up with the assignments. You keep up with the assignments, so you will do well.

The first premise says that if you keep up, you'll do well. The second premise says that you do keep up. You have to supply the conclusion, which is the claim that you will do well. The argument affirms the antecedent, so it is valid.

6. My car only dies when the temperature is below freezing. But it's below freezing today, so it will die.

Hint: what is the necessary condition here?

7. If Smith embezzled the money, then Jones was involved in the crime. But Smith didn't embezzle it. So Jones wasn't involved.

Denying the antecedent. Invalid.

8. If Wilbur is an uncle, he couldn't have been an only child. And he is an uncle.

You must supply a missing conclusion here; then it should be easy.

9. If wishes were horses, then beggars would ride. But beggars don't ride. So it looks like wishes aren't horses.

Denying the consequent. Valid

10. If everyone here at OU were doing well with the current course requirements, there would be no need to change the requirements. But some people are not doing well.

Unstated conclusion: we should change the present course requirements. Denying the antecedent. Invalid.

11. If OU has a winning record in the Big 12, then if all of their players are healthy, they will do well in the tournament. And they have a winning record.

Affirming the antecedent. Valid.

12. If Tom's prints are on the gun, then he is guilty. So he must be innocent, because those weren't his prints on the weapon.

Denies the antecedent. Invalid (being innocent is the opposite of being guilty).

### 3. Exercises on Necessary and Sufficient Conditions

1. "You will not pass this course unless you pass the final." Here the claim is that passing the final is a necessary condition for passing the course.
2. "John will marry Sue only if she agrees to have three children." Here the claim is that agreeing to have three children is a necessary condition for his agreeing to marry her.
3. Vixens are female foxes. Being a vixen is both a necessary and a sufficient condition for being a female fox. It is a definition of 'vixen', and definitions typically involve both sorts of conditions.
4. "You'll won't be happy, if you buy it at Sturdleys." Here the claim is that buying it at Sturdley's is a sufficient condition for being unhappy.
5. "A person is a brother just in case he is a male sibling." See the answer to the question on vixens above.

## **Part III**

# **The Acquisition and Retention of Information**



### **Part III. The Acquisition and Retention of Information**

Reasoning has to begin with something. Arguments require premises. Sometimes the *inputs* (i.e., the premises) for our reasoning are conclusions drawn in earlier inferences, but in the end most of our knowledge can be traced back to two sources: observation and the claims of other people. In this part we will study perception, other people as sources of information, the internet, memory, and the effects of emotions on reasoning.



## Chapter 4

# Perception: Expectation and Inference

**Overview:** Perception may seem unrelated to reasoning; it's natural to think that if we simply turn our heads in the right direction, we'll take in whatever is there. But in fact perception involves much more than passively receiving incoming information; perception is something we *do*. Indeed, perception is very much like inference, and it often goes awry because of the same factors—context, expectations, biases, wishful thinking—that lead to flawed reasoning.

### Contents

---

<b>4.1</b>	<b>Perception and Reasoning</b> . . . . .	<b>66</b>
<b>4.2</b>	<b>Perception is Selective</b> . . . . .	<b>66</b>
<b>4.3</b>	<b>There's More to Seeing than Meets the Eye</b> . . . . .	<b>67</b>
	4.3.1 Information Processing . . . . .	68
<b>4.4</b>	<b>Going Beyond the Information Given</b> . . . . .	<b>68</b>
<b>4.5</b>	<b>Perception and Inference</b> . . . . .	<b>70</b>
<b>4.6</b>	<b>What Ambiguous Figures Teach Us</b> . . . . .	<b>71</b>
<b>4.7</b>	<b>Perceptual Set: the Role of Expectations</b> . . . . .	<b>73</b>
	4.7.1 Classification and Set . . . . .	75
	4.7.2 Real-life Examples . . . . .	76
<b>4.8</b>	<b>There's more to Hearing, Feeling, . . .</b> . . . . .	<b>77</b>
	4.8.1 Hearing . . . . .	77
	4.8.2 Feelings . . . . .	78

---

<b>4.9</b>	<b>Seeing What We Want to See</b>	<b>79</b>
4.9.1	Perception as Inference	81
4.9.2	Seeing Shouldn't be Believing	81
<b>4.10</b>	<b>Chapter Exercises</b>	<b>83</b>

---

## 4.1 Perception and Reasoning

Perception is related to reasoning in several important ways.

1. Our reasoning is often based on premises that describe what we see or hear. Furthermore, such premises are usually thought to be especially secure and trustworthy.
2. Perception requires us to go beyond the information given to us by the surrounding environment. This leap beyond the incoming information involves something very much like reasoning or inference.
3. This perceptual inference can be influenced by the context, our expectations, and even our biases, desires, and self-interest. These are the very same things that often lead to faulty reasoning.
4. Since perception is susceptible to various sorts of errors, we need critical reasoning to evaluate claims about what we (and others) perceive.

Reasoning has to begin with something, and we can trace many of our beliefs back to perception, to information we acquired from our environment. Perception is the interface between the mind and the world. So by starting with perception we begin at the beginning. But we will also find that many of the things we learn about perception apply, with modest changes, to many aspects of reasoning.

## 4.2 Perception is Selective

### Filtering and Selection

There is a lot going on in the world around us. But we aren't overwhelmed by sensory overload, because perception is *selective*. Some of the selection and filtering occurs at the neural ("hardware") level. For example, the human visual system is only sensitive to a small band of the electromagnetic spectrum in the interval between ultraviolet and infrared electromagnetic radiation. We can't see information conveyed by X-rays (like the people in science fiction stories who can see through



walls) or by infrared light (unless we wear special goggles, like those that allow us to make out shapes at night). Similarly, we can't hear the high-pitched sounds dogs can hear or make sense of the intricate pattern of shrieks emitted by bats to navigate in the dark.

### Attention: Further Filtering

Some perceptual selectivity is "wired in," but some also depends on our beliefs, emotions, desires, and other things that aren't part of our in-born physiology. When you enter a large room full of noisy people, it is just a loud din. But once you strike up an interesting conversation with someone, you tune out most of the noise around you and focus on them. You have no idea what people across the room are saying—until one of them mentions your name. Then, all of a sudden, their words leap out at you.

Experiments show that this sort of phenomenon is common. We filter out a lot of information, but when it becomes relevant, we often tune it back in. Some of the factors that lead us to focus on some things while ignoring others—things like expectations and emotions and desires—bear on our topic of reasoning, but here we will focus on an even more direct connection.

## 4.3 There's More to Seeing than Meets the Eye

When we talk about the input to the visual system we naturally think about the eyes. Light passes in through the lens of each eye and is projected onto rods and cones at the back of the retina. But our perception is much different and much richer than this input. The lens reverses things from left to right and turns everything upside down. The image on the back of the retina is not homogeneous, but punctate (because the rods and the cones are separate, discrete units). Furthermore, the eye is constantly moving around, and we blink frequently.

In short, the image on the back of our retina is two-dimensional, upside down, punctate, and it jumps all over the place. But we see a three-dimensional world, right side up, full of familiar objects that don't constantly leap about.

A related point can be made about *perceptual constancies*, such as size, shape and distance constancies. Hold your hands up in front of your face, with the left hand close to your eyes and your right hand as far away as possible. Do your hands look the same size? Walk over to the wastebasket, then back away while keeping it in view. Does it seem to become smaller as you back away from it? Circle around and view it from different perspectives. Does it seem to change shape as you move? The fact that the wastebasket seems the same size as you back away from it (so that

the smaller retinal image of it is smaller) is known as *size constancy*. The fact that it doesn't seem to change shape as you circle around it (so that the retinal images of it do change shape) is known as *shape constancy*.

As long as you don't get too far away the size and shape of the wastebasket seem to remain constant; your perception of it is the same. Yet its image on your retina becomes smaller as you back away from it, and the image changes shape as you change your orientation with respect to the wastebasket. Here the *input is different*, but the appearance of the object *remains the same*. So something is going on that involves more than just the images on your retinas. There is some debate over how to explain size constancy, but there is some evidence that in part it is caused by the perceptual system's enrichment of the sensory input.

### 4.3.1 Information Processing

**Information processing:**  
transforming information in  
the input into information in  
the output

In the cases considered thus far the perceptual system seems to *enrich* the visual input. It does not just passively register it, it *does something* to it. It is often useful to think about such processes in terms of information processing. Suppose that you program your computer to accept certain *inputs* (say a series of number between 0 and 100), and then have it manipulate or process them so that it generates a certain *output* (say the average values of the numbers in the input).

Many cognitive processes also involve information processing. In the case of visual perception, the *input* consists of the stimulation of rods and cones at the back of the retina, the relative orientation of our eyes to one another and perhaps various factors involving the orientation of our body. And the *output* is the perceptual experience we have when we see something.

**Bottom-up processing:**  
transmission of information  
from sensory receptors up to  
the brain

In the cases considered so far the perceptual system enriches the visual input in ways that do not involve our beliefs or desires. This is called **bottom-up processing**. The idea here is that our nervous system records sensory stimulations and passes the information on *up* to the brain. But we see with our minds as much as with our eyes; we also engage in "top-down" processing that involves something very much like inference. So, before turning to it, it will be useful to recall a few facts about inference.

## 4.4 Going Beyond the Information Given

The word 'inference' is almost synonymous with 'reasoning'. We draw an inference when we begin with one or more beliefs (our premises) and use them to arrive at a conclusion. We start with a body of information (or perhaps misinformation, or some of each) and arrive at a new piece of information. Suppose that we know that

94% of the people in the critical reasoning class are from Oklahoma and later we learn that Tom is in the course. We infer that Tom is (probably) from Oklahoma. We arrive at this conclusion on the basis of the two earlier pieces of information.

Many inferences are drawn so quickly and automatically that we don't even notice it. Consider a quarterback running an option play. As he runs along the line he must read the location and trajectory of the on-coming defensive players. Then, depending on what he sees, he decides whether to pitch the ball or to keep it. A good option quarterback will usually gather and assess the relevant information very quickly and make the right decision in a split second.

When we first learn a skill (like riding a bike or skiing or typing) it seems awkward and unnatural. We have to rehearse every step in our minds as we struggle to get the hang of things. But as we become more adept, we no longer have to think about each step consciously; indeed, it often becomes difficult to even say what it is that we do (many good touch typists find it difficult to describe the locations of various keys).

There is some very recent evidence that once we acquire a skill, the ability to use it is relocated in a part of the brain that isn't accessible to consciousness. But the important point here is that many of the skills we acquire involve gathering and assessing information and drawing inferences from it.

Scientists still have much to learn about cognition, but at least this much is clear: much inference takes place very rapidly and below the threshold of consciousness. Once we grant this, it won't seem so odd to think that perception often (perhaps always) involves something very similar to inference. In fact, it is so similar that some scientists think of perception as a special kind of inference.




Figure 4.1: Beyond the Information Given

The conclusion of a deductively valid argument doesn't contain any information that wasn't already contained (often in a far from obvious way) in the premises. By contrast, the conclusion of an inductively strong argument does contain new in-

formation; its conclusion *goes beyond* the information given in its premises. An inductively strong argument involves a jump, an *inductive leap*, to new information in a way that a deductively valid argument does not. This is why it is possible to have an inductively strong argument with true premises but a false conclusion.

## 4.5 Perception and Inference




Figure 4.2: Necker Cube

Perception is a lot like inference that makes an inductive leap. There is a huge amount of information around us, but some of the information we need (or want) is not present in the input to the visual system. We will see why this works in the case of visual perception, then we will briefly note how similar points apply to some of the other senses.

The object in Figure 4.2 is known as a *Necker cube*. We can see it *as* a cube with a given face (ABCD) or, when it reverses, *as* a cube with a quite different face (EFGH).

Figure 4.3 on the next page shows more dramatic examples of *reversing figures*. The physical input to the visual system is the same whether we see the stairs (in the left subfigure) on the bottom *or* on the top. So what *accounts for* the difference when we see it first one way, then the other?

What do you see in Figure 4.4 on the facing page? It can be seen either as two faces peering at each other or as a vase. Figures like this, which can be seen in more than one way, are called *ambiguous figures*.

Now consider Figure 4.5 on page 72. What sort of person have we here? Like our earlier figures, she can also be seen in either of two ways. Look at the figure before you read on (she can be seen as either a young woman or an old woman; the chin line of the young woman is part of the nose of the old woman).




Figure 4.3: Other Reversing Figures



Figure 4.4: Faces or Vase?

Next consider the creature in Figure 4.6 on the following page. You can see it as a duck (looking off to the left) *or* as a rabbit (looking off to the right).

#### 4.6 What Ambiguous Figures Show: Expectations and Set

Ambiguous figures are intriguing. They aren't typical of the things we normally see, though, so why spend so much time on them? The answer is that they tell us something very important about the *causes* of our visual experiences.

Let's begin with a simple example of causation. Suppose that you enter the same sequence of numbers on your computer which you have programmed to cal-



Figure 4.5: An Ambiguous Woman



Figure 4.6: An Ambiguous Animal

culate averages. Then you are asked to enter the same numbers on my computer, which I tried to program to work the same way that yours does. But the two computers come up with different answers. How might we explain this?

One hypothesis is that you mistakenly entered different numbers in the two computers. But you double check and find that you gave the same input to each computer. Now what could explain the different results? Since the input is the same—it is *held constant*—in the two cases, the difference must have something to do with what goes on inside each computer.

So far, so good. But what is it in the computers that accounts for the different outputs? Perhaps one of the programs has a bug in it. You could test to see whether this was the case by working through each program. If they are the same, then

the difference in the two computers' behavior must have something to do with the hardware or the other programs in the computer. (At this point you might want to call in your friendly hacker, Wilbur.)

We often reason about causation in this way. In a later chapter we will examine the intricacies of such reasoning, but here it is enough to note that ambiguous figures allow us to learn something about the causes of perception in the same way that we learned something about the causes of the different outputs of the two computers.

When you look at the Necker cube or any of the other ambiguous figures, the *input* to your visual system is the same no matter how you perceive the figure. It is the same when you see the two faces as it is when you see the vase. Yet your visual experience is different.

Since the *input is the same* in the two cases, the difference must have something to do with what happens inside you *after* the image is formed. It involves the way the input is processed. Moreover, if we can manage to hold further factors constant, we may be able to zero in on the factors that affect the internal processing.

In short, perceptual constancies suggest that we can have *different inputs* while having the *same output* (the same visual experience). This suggests that sameness of input is *not a necessary condition* for having the same experience. And our ambiguous figures show that we can have *the same input* while having *different output* (different visual experiences). This shows that sameness of input is *not a sufficient condition* for having the same experience.

1. Perceptual constancies suggest that having the same sensory input is not necessary for seeing the same thing.
2. Ambiguous figures show that having the same sensory input is not sufficient for seeing the same thing.

The moral is that the mind is active. But what determines the nature of the active role it plays? The next few examples give us part of the story.

## 4.7 Perceptual Set: the Role of Expectations

What is the character in the middle of Figure 4.7 on the next page? If we block out the *12* on the left and the *14* on the right we just see the column, and we naturally, almost automatically, see the middle character *as a B*. But block out the *A* and the *C*, and it instead looks like a *13*.

Here the *context* influences our expectations, so that we tend to see what the context leads us to expect we would see. To test this claim we could show one



Figure 4.7: What's that in the Middle?

group of people the three characters on the row and another group the three characters in the column. When we do this we find that context does indeed influence how people see things. This example also shows how *cultural issues*, our having the language and alphanumeric characters that we do, affect how we see things. People from a much different culture might not see this figure as anything but a series of meaningless marks. In fact there is evidence that the perception of certain visual illusions varies from one culture to another.

Can you think of a context that might lead people to see some of the ambiguous figures in one way rather than another? Figure 4.8 should give us a clue. Here we plop our ambiguous creature down into two different contexts.



Figure 4.8: Two Group Portraits

Now answer the following questions:

1. If you saw the group portrait of the ducks first, how do you think you would have interpreted the picture of the lone creature?
2. What if you had instead seen the group portrait of rabbits?
3. Do you think it would have made any difference if someone had first said to



you "I'm going to show you a drawing of a duck?"

4. How do you think that people who had seen lots of rabbits but no ducks would see it?

These are *empirical questions*, and the only way to answer them is to check and see what people actually do under such conditions. But we know the effect these examples have on us, so we have something to go on. In a moment we will consider some further cases that will provide more evidence about the correct answers.

What you *expect* to see can strongly influence what you *do* see. Psychologists say that your beliefs and expectations (and, as we will see in a bit) your desires constitute your **perceptual set**.

**Perceptual set:** how we are primed to see things in a given setting

The context helps determine your perceptual set, because it influences what you expect to see. In some contexts you expect to see one sort of thing, say a duck. In other contexts you may expect to see something else, say a rabbit. In normal situations you would be upset if coins simply began to vanish into thin air. But when you watch a magician, you expect coins that seem to disappear.

#### 4.7.1 It just isn't in the Cards: Classification and Set

Context is one important determinant of perceptual set, but it isn't the only one. An experiment done in 1949 supplies some surprising insights into the way that our expectations are influenced by our ways of classifying things. Jerome Bruner and Leo Postman presented subjects with a series of five-card hands of playing cards that flashed briefly before them for periods of a second or less (nowadays this would be done on a computer screen). The hands contained many normal cards, but some cards were anomalous. On these cards the hearts were black and the spades were red.



It took people longer to recognize these trick cards than to recognize normal cards. At one or more points in the experiment almost all of the subjects reported an anomalous card as being normal. For example, they assimilated the black three of hearts to a normal three of hearts (here the shape dominated) or to a normal three of spades (here the color dominated).

The subjects had a system of classification: cards come in four suits (hearts, diamonds, spades, and clubs) with cards in the first two suits being red and cards in the last two suits being black. This led them to expect certain things, and in some cases these expectations led them to misperceive a trick card.

The influence of our expectations and beliefs and desires on perception involves **top-down processing**. Both bottom-up and top-down processing are important. There is still some controversy about the relative importance of each (and of additional sorts of processing we won't go into here); it might be thought, for

**Top-down processing:** processing of visual information that involves expectations and beliefs

example, that top-down processing only plays a substantial role in cases where the object of perception is ambiguous. But even if some scientists overestimate the significance of top-down processing, the following examples show that beliefs and expectations, desires and emotions, can have a substantial impact on how we see and interpret normal objects outside of the laboratory.

#### **4.7.2 Real-life Examples**

There is sometimes disagreement about how to interpret laboratory studies, and even when there isn't they can seem a bit artificial. But there can be little doubt that our perceptual sets influence the ways that we see things in the world outside the laboratory—sometimes with disastrous results.

##### **It looked like a Bear**

Elizabeth Loftus and Katherine Ketcham report a case of two men who were hunting bear in a rural area of Montana. After an exhausting day in the woods night was falling and they were making the trek home. On the way they talked about bears. Suddenly, as they rounded a bend, a large, moving object loomed up ahead of them. Both men took it for a bear, raised their rifles, and fired.

It turned out to be a large, yellow tent with its flap blowing in the wind. A couple was inside, and the woman was killed. The hunter who shot her was tried (and convicted) of negligent homicide. The jury found it incomprehensible that he could have mistaken a yellow tent for a bear. But he was primed to see a bear; he had bears on the brain. His perceptual set played a powerful role in leading him to see things in the way that he did. Most of us have probably made similar mistakes, though fortunately with less tragic results.

##### **Schizophrenic or Normal?**

In 1973 researchers had a healthy, normal adult check himself into a mental institution with complaints of hearing voices. He was classified as a schizophrenic and admitted. Once in the hospital he never said anything about the voices, and he behaved like the normal adult that he was. But the members of the hospital staff had been led to expect a schizophrenic, and that is how they continued to see him.

They recorded his “unusual” behavior (he spent a lot of time writing down his observations), often talked to each other as though he wasn't present (this is common behavior in psychiatric wards), and didn't realize that he was behaving normally. Once they had a label, they didn't check to see how well it fit. It took, on average, almost twenty days for the subject to get himself released. This is not a

case where people misperceived something they only glimpsed for an instant. Here the staff was around the subject for almost three weeks. Interestingly, some of the hospital patients were much quicker to see that the subject was normal (what might explain this?).

### Everyday Errors

Misperception is far from rare. Here are three examples; can you think of some more?

- Most of us have seen a person in the distance and thought they were someone we know. When they got closer, we realized that they don't look anything like the person we expected.
- Proof-reading provides similar examples. As you read back over the first draft you *expect* to see words spelled in a certain way, and things often look as you think they should. It is very easy to read right past the 'teh' there on the page—though once you do notice it, it jumps off the page at you.
- If you don't normally live alone and find yourself home all by yourself (especially if you have just seen a scary movie), shadows in the backyard and sounds in the attic can take on new and sinister forms.

## 4.8 There's more to Hearing, Feeling, ...

### 4.8.1 Hearing

If you hear people speaking a language you don't understand, you are unlikely to perceive breaks between many of their words; you won't know which sounds are single words and which aren't. Listen carefully to an English sentence spoken at a normal speed, and you'll realize how the sounds run together. But someone who knows the language and expects to hear normal words will perceive discrete words rather than just one long run-together sound.

Context also affects how we hear things, particularly words. The phrase 'eye scream' is interpreted differently in the sentences 'I scream when someone jumps out and surprises me' and 'I love rocky road ice cream'.

Back in the 1960s the rock-and-roll band the Kingsmen had a smash hit with "Louie Louie." It wasn't easy to discern the words, and there were various accounts of what they were. The set on the sheet music was harmless enough, but another set, circulating widely in the teenage underground, would have kept the song off the air. It turned out that if you gave people one set of words to read before hearing

the Kingsmen's rendition, they would hear those words. But if you gave them the other set, they would hear those.

### The Phoneme Restoration Effect

When people hear the following sentence

It was found that the \*eel was on the \_\_\_\_\_.

where \* is a missing sound, they automatically fill the \* in so that they think that they hear a normal English word. The way that they fill it in depends on what word is placed in the blank. When the following words were put into the blank underlined space

1. axle
2. shoe
3. orange
4. table

it determines what people think they hear. For example, they think they hear the word 'wheel' (in the axle case). What words did they think they heard in the other cases?

The \* represents what linguists call a 'phoneme', and so this phenomenon is known as the *phoneme restoration effect*. It is particularly interesting because the relevant part of context (the word inserted in the blank) is yet to come when subjects filled in the missing phoneme.

Similar points apply to other sensory modalities. Suppose that you are squeamish about bugs. You go on a camping trip and around the campfire people swap stories about scary insects, people they know who died from spider bites, and the like. That night as you are drifting off to sleep a blade of grass brushes against your cheek—but when you first feel it, it probably won't feel like a harmless blade of grass. Or you may be enjoying a tasty meal—until you learn it contains some ingredient you don't like or find disgusting (that hamburger tasted wonderful—until you learned that it was horse meat).

## 4.8.2 Feelings

### Physiological States and Context

Expectations and beliefs can even influence our physiological state. In a study done in the 1970s two psychologists had male undergraduates take a drink. Some of the drinks contained alcohol and some didn't. Then some people in each group

were told that their drink contained alcohol and some were not told this. Finally, a female assistant then walked in, sat down, looked the subject right in the eye, and begin talking to him—which made many of the subjects nervous.

We all know that nervousness affects heart rate. It turned out that subjects who thought they had been given a vodka tonic showed smaller increases in heart rate than subjects who thought they'd had only a glass of tonic. Whether the subjects *really* had been given alcohol didn't affect their heart rate. But whether they *thought* they had been given alcohol did.

Expectations and context can exert a powerful effect on our perceptions and feeling. In some situations they lead to a placebo effect. The **placebo effect** occurs when people are given a pill or a shot consisting of chemicals that won't affect their illness or disease. If they *think* that it is genuine medicine, they often get better, even though they only took a sugar pill. In a later module we will see that this effect is so powerful that experiments on the effectiveness of drugs must be designed to guard against it.

## 4.9 Seeing What We Want to See

Perhaps expectations can affect our perceptions, but can our desires and emotions (which account for so much fallacious reasoning) have an impact on them? There is strong evidence that they can.

### The Football Game

In a classic study from the 1950s Albert Hastorf and Hadley Cantril examined biases and their effect on perception. In 1951 Dartmouth and Princeton met on the football field. The game was unusually rough, and there were several injuries and many penalties on both sides. After the game, partisans of both teams were upset.

When Hastorf and Cantril asked two groups of students, one from each university, which team started the dirty play, the groups from the two universities gave quite different answers. Of course they may have heard about the game from someone else, so to study the effects of actually *watching* the game, Hastorf and Cantril asked a group of boosters of each school to watch a film of the game and record each penalty they noticed. Princeton boosters saw many more Dartmouth penalties than Dartmouth boosters did. Here again expectations influenced perception. But in this case peoples' expectations were influenced by which school they identified with.

**Hostile media**

**phenomenon:** most people who think the media is biased think it's biased *against* their views

**The Biased Media**

About ten years ago several psychologists studied the way that voters viewed the media. It turned out that about a third of the respondents thought that the media had been biased in their coverage of Presidential candidates. There is nothing too surprising about this, but in 90% of the cases where people discerned a bias, they perceived it as a bias *against* their candidate. This has become known as the *hostile media phenomenon*. Psychologists found this phenomenon regardless of the candidate involved. They also found similar outcomes when the issue was media bias in the presentation of other sorts of news events. Here one's values and desires play a role in what one sees or, at least, in how one interprets it.

Such things also happen closer to home. Most of us are prone to see "bad" officiating calls when they go against our team, but we don't notice many that go against the opponent. After a game people often complain that their team lost because of poor officiating, but few say that bad officiating gave their team the victory. One way to see the influence of this bias is to try to imagine how the officiating calls in a game would be viewed by one of the opposing team's fans.

The influence of our desires on perception isn't limited to the sports world. Many parents are unable to see what their children are doing (e.g, abusing drugs) because they can't bring themselves to believe their child would do that. People in a relationship may be unable to see obvious flaws in the person they care about. Of course not all biases lead us to think the best of some one else. If Wilbur is prone to jealousy, harmless and friendly behavior on his wife's part may look like flirting to him.

**Person Perception**

Our perceptions of other people are influenced by our perceptual set just as much as any of our other perceptions are. For example, our set may be influenced by stereotypes and biases that lead us to expect to see certain things, and sometimes this can actually lead us to see them in that way. We also have stereotypes about people who dress in certain ways, sport particular hair styles, have certain physiological characteristics, and so on, and these also influence our perceptual set.

It is natural to wonder about the effects of over-simplified classifications, expectations fostered by parents or peer group, and biases and desires might affect our perception of people of different races or from opposing political groups. The topic is so important that we will reserve an entire chapter for it later in the course. But it is important to note now that the things we have learned about in this chapter are not just about ambiguous figures. They turn up in all sorts of situations, including the social situations that matter most to us.

### 4.9.1 Perception as Inference

It is difficult to escape the conclusion that perception works a lot like inference that goes beyond the information that we have. In fact one school of thinkers, beginning with the German Hermann von Helmholtz in 1866, holds that perception *is* a species of inference. But for our purposes it is enough to realize that in one very important way perception is like inference. The input from the outside world, consisting of light rays and probably some less obvious things, is analogous to the *premises* of the inference. And the actual perceptual state we experience is analogous to the *conclusion*.

### 4.9.2 Seeing Shouldn't be Believing

We will see over and over how biases and self-interest and wishful thinking lead to fallacious reasoning. And the fact that they can influence *what* we see, or at least *how* we see it, suggests that perception can be flawed for many of the same reasons that reasoning can. This is a serious problem, because we have a very strong tendency to think that our perception is accurate. Indeed, we even tend to put a lot of faith in what other people claim to see (eye witness testimony carries great weight in the courtroom). But errors are very possible here, and so we often need to subject our perceptual beliefs to scrutiny.

We will find that many of the things we have learned about perception turn up repeatedly in our study of reasoning. Here is a list of some of the key points we will meet on future occasions.

1. It is important to us to make sense of the world around us, to explain what happens and to fit it into a coherent and organized pattern. In perception we strive to make sense of the things we see and hear. Memory and inference involve similar attempts to make sense of things.
2. Perception, memory, and inference are strongly affected by several factors, and often these things lead to errors. These factors include:

1. Context
2. Our beliefs and expectations
3. Our wishes and desires.

3. Our perceptions and reasoning can be influenced, even distorted, by these factors, but there are limits to their influence. If our beliefs and desires did completely determine what we saw, we wouldn't be able to function on our own for even a day. Not all visual illusions involve ambiguous figures, and some of them actually demonstrate the *limitations* of perceptual set. Figure 4.9 on the next page is known

as the Müller-Lyer illusion. The line with the out-going fins looks longer than the line with the in-going fins, but if you measure them you will find that they are the same length. Even once you know this, however, your belief that they are the same length and even a strong desire to see them as having the same length are not enough to enable you to see them as having the same length.




Figure 4.9: Müller-Lyer Arrow Illusion

But although there are limits on how wrong we can be, we often do make mistakes, even in situations that matter greatly to us. Knowing about these pitfalls in perception is a first step in guarding against such errors.<sup>1</sup>

<sup>1</sup>The phrase *going beyond the information given* was used by Jerome Bruner and his collaborators to describe various types of inductive inference. It is, in fact, an excellent label for *all* types of inductive inference. You can find a discussion of perceptual constancies in any introductory text in psychology. The influence of culture on the perception of the Ponzo illusion can be found in H. W. Leibowitz, et. al, "Ponzo Perspective Illusion as a Manifestation of Space Perception," *Science* 166 (1969); 1174–1176. The classic card experiment described above was conducted by Bruner and Leo Postman, "On the Perception on Incongruity: A Paradigm," *Journal of Personality*, 19 (1949); 206–223. The bear hunting case is described by Elizabeth Loftus and Katherine Ketcham, *Witness for the Defense*, St. Martins Press: 1991, p. 22-23. The normal patient in the mental hospital is discussed in David Rosenhan's "On Being Sane in Insane Places," *Science* 173 (1973); 250–258. Schachter describes some of his research in "The Interaction of Cognitive and Physiological Determinants of Emotional State," in L. Berkowitz, ed., *Advances in Experimental Social Psychology*, Vol I; Academic Press, 1964. The alcohol study is reported in G. T. Wilson and D. Abrams, "Effects of Alcohol on Social Anxiety and Physiological Arousal: Cognitive versus Pharmacological Processes," *Cognitive Research and Therapy*, 1, 1977; 195–210. The conflicting interpretations of the Dartmouth-Princeton football game are described in A. R. Hastorf and H. Cantril, "They Saw a Game: A Case Study," *Journal of Abnormal and Social Psychology*, 49, 1954; 129–134. The studies on the hostile media are reported in R. P. Vallone, L. Ross, and M. R. Lepper, "The Hostile Media Phenomenon: Biased Perception and the Perceptions of Media Bias in Coverage of the Beirut Massacre," *Journal of Personality and Social Psychology*, 49 (1985); 577–585. A spirited defense of the view that perception is very much like inference may be found in Richard Gregory's *Eye and Brain: The Psychology of Seeing*, 4th ed. Princeton University Press: 1990. Gibson's alternative view is developed in his *The Ecological Approach to Visual Perception*, Houghton Mifflin: 1979. Jerry Fodor discusses the way some illusions show the limitations of perceptual set in "Observation Reconsidered," *Philosophy of Science* 51 (1984); 23–42.



## 4.10 Chapter Exercises

1. Think of a case where you were certain that you saw something or someone, but where on closer examination, you discovered that you hadn't really seen that person or thing at all (or at least that the thing or person you saw looked very different from the way that you first thought they did). Write a paragraph describing this situation; include a discussion of things that might have led to the misperception.

2. We have some tendency to selectively perceive what we expect and hope to see. Describe, and comment on, an example in which you (or a person close to you) have done this. What cognitive or motivational factors were at work in your perception?

3. Context can influence our expectations and so it can influence our perceptual set. Describe one way in which you might set up a context in which you think people would be more likely to see the vase rather than the face. Now describe a context where people might be more likely to see the face. How would you test your hypotheses about this?

4. Context can influence our expectations and so it can influence our perceptual set. Describe one way in which you might set up a context in which you think people would be more likely to see the old woman rather than the young woman. Now describe a context where people might be more likely to see the young woman. How would you test your hypotheses about this?

5. The staffs of mental hospitals took, on average, three weeks to discover that the normal person who had been admitted as a schizophrenic was in fact normal. The patients were much quicker to see that he was normal. What do you think explains this difference?

6. Describe a case where your expectations or desires or the context you were in seems to have led you to interpret something you felt (with your skin, your tactile sense) one way that, with different expectations or set, might have been interpreted another way.

1. What do you think caused you to interpret it the way that you did?
2. How could you test the hypothesis you constructed to answer the previous question?

Give similar examples involving taste and smell.

7. Describe a case where your expectations or desires or the context you were in seems to have led you to interpret your emotions or mood one way that, with different expectations or set, might have been interpreted another way.

1. What do you think caused you to interpret it the way that you did?
2. How could you test the hypothesis you constructed to answer the previous question?

## Chapter 5

# Evaluating Sources of Information

**Overview:** Most of our knowledge and reasoning is based on things we learn from other people. In this chapter we will focus on other people as sources of information. Information acquired in this way is often reliable, but it isn't foolproof; people make mistakes, and sometimes they intentionally misrepresent things. So we need ways to decide when it is reasonable to accept their claims and when it is better not to. In this chapter we will examine various sources of information and develop some guidelines for separating reliable sources of information from unreliable ones.

### Contents

---

<b>5.1</b>	<b>Other People as Sources of Information</b> . . . . .	<b>86</b>
5.1.1	Information: We need something to Reason About . . .	86
<b>5.2</b>	<b>Expertise</b> . . . . .	<b>88</b>
5.2.1	What is an Expert? . . . . .	88
5.2.2	Fields of Expertise . . . . .	89
<b>5.3</b>	<b>Evaluating Claims to Expertise</b> . . . . .	<b>89</b>
<b>5.4</b>	<b>Who Do we Listen To?</b> . . . . .	<b>93</b>
5.4.1	Faking Expertise: The Aura of Authority . . . . .	94
5.4.2	Appearing to Go Against Self-interest . . . . .	96
<b>5.5</b>	<b>Evaluating Testimony in General</b> . . . . .	<b>97</b>
<b>5.6</b>	<b>Safeguards</b> . . . . .	<b>99</b>

---

 5.7 Chapter Exercises . . . . . 102
 

---

## 5.1 Other People as Sources of Information

In this chapter we will focus on other **people as sources of information**. People's claims are sometimes called *testimony*, but don't let this word mislead you. We will use it to cover any information (and misinformation) we acquire from other people. In this sense testimony includes the things we hear directly from others, read in newspapers or books, see on television, find on web sites, and so on.

Think about how much of your knowledge is acquired from others. Your friends tell you about things they have seen or heard. You learn new things when you read the morning paper or a chapter in a textbook or an email message from someone far away. You acquire information when you watch the news or surf the net.

A large society like ours consists of a vast social network in which everyone relies on others for information, and written records and oral traditions extend this network back into the past. If you suddenly forgot all of the things you learned from others, there wouldn't be much left. You wouldn't even have a language in which you could ask yourself how much you knew.

In today's rapidly changing world much of what you learn in college will become outdated rather quickly. Many of your grandparents, and perhaps even your parents, had just one or two jobs during their adult life. But the swift pace of globalization and technological innovation make it likely that you will have a succession of jobs once you graduate. Hence it is important for you to *learn how to learn*, and a key part of this is learning how to acquire and evaluate information.

### 5.1.1 Information: We need something to Reason About

Testimony bears on reasoning for two closely related reasons.

1. The *premises* of our reasoning are very often based on things we learn from others.
2. We frequently need *background information* to know whether the premises of an argument are plausible or whether they omit relevant information, and often we must rely on others to supply this information.

**Our premises are often based on the claims of others.** Many of our arguments rely on premises that we get from others. Suppose, in an effort to create

jobs, your City Council proposes that a nuclear power plant or a toxic waste dump be built near your community. Most of us would care enough to argue for our position on this issue. But we would have to rely on the claims of experts about the risks and reliability of such facilities to support our conclusions about whether a potentially dangerous plant should be built.

**Evaluation of arguments often requires background knowledge obtained from testimony.** There are three critical questions to ask about any argument we encounter. Do its premises support its conclusion? Are its premises plausible? Has any relevant information been omitted? Often we cannot answer the last two questions without having a good deal of background knowledge, and frequently we must rely on other people to supply it.

Here's an example of an issue you will face as a parent. The average child will see over 8000 depictions of murders on television before they graduate from grade school. Suppose your neighbor urges you to sell your TV so that your little Wilbur won't grow up to be an axe murderer. In trying to evaluate the argument that TV violence is dangerous for your children you need background information about the lasting effects of seeing frequent violence on television. Will seeing violent acts on TV make Wilbur more violent.

Or suppose someone argues that we should retain capital punishment because it deters murders. Here you need background knowledge to know whether their premise that capital punishment really does deter murders is true. And you also need to know whether they have omitted information; for example, it may seem to deter murders in one state (which they mention), but not in other states (which they don't).

In a heterogeneous, highly-technological society like ours *experts* are an especially important source of information. We rely on our dentist when we break a tooth, on the telephone repairperson when our phone goes dead, on the morning paper to tell us how OU did in last night's game, on *Consumer Reports* for help in buying a used car.

It requires *effort* to obtain useful information. If we sit around passively and hope that useful information will come our way, we will be in trouble. But effort alone isn't enough. It is often difficult to distinguish genuine experts from self-styled experts who don't really know what they are talking about. Furthermore, even genuine experts are sometimes biased, and in many cases the experts disagree with one another. Finally, on some matters there may not be any experts at all. This means that we need techniques for identifying experts, and we also need to know what to do when they can't be found or when they disagree.

We will begin by examining experts in recognized fields like medicine and law and professional football. We will then consider several things that make people appear credible; the goal here will be to discover ways in which people can ap-

pear to be experts when they really are not. We will then turn to testimony more generally. We will find that the issues that arise in evaluating the claims of experts are similar to those we should use for evaluating *any* sort of testimony, whether it comes from a top expert in some field or a stranger on the street.

## 5.2 Expertise

The **identification of experts** involves finding experts on the relevant topic. This may require us to separate genuine experts from the frauds, quacks and charlatans. We often have no alternative but to rely on experts, but in the end it is still up to each of us to decide *when* we will rely on experts, *which* experts we will rely on, and *what to do* when the experts disagree. In the end we have to decide for ourselves, and in matters of importance it is best to base our decision on the best information we can get.

### 5.2.1 What is an Expert?

An **expert** or **authority** is someone who knows a lot about a particular field. We will use the words ‘expert’ and ‘authority’ in a broad sense, so that we can consider a wide variety of **sources of information** as potential authorities. So when we talk about an authority, we don’t mean an authority figure who is in charge of other people, but someone (or something) who is an authority in their field. In this sense authorities include individual people, newspapers, textbooks, encyclopedias, TV programs, web sites, think tanks, and so on.

Sometimes expertise may be embodied in a skill that is difficult to describe. A doctor who has seen many patients with a certain disease may be able to recognize it even if she has trouble saying precisely how she does so. Still, if she is usually right, she is an expert at recognizing the disease. There are two important facts about experts.

*Experts needn’t be infallible*—who is? If this were a requirement, then there wouldn’t be any experts. But an expert will still be a better source of information than someone who isn’t an expert. Indeed, if the experts were right about something important 60% of the time whereas nonexperts were only right 50% of the time, you would still be better off relying on an expert.

*Expertise comes in degrees.* Most dentists are experts on teeth. But a conscientious dentist who has practiced for twenty years is likely to know more about molars and bicuspid than a lackadaisical dentist who has been practicing for three days. But when you break a molar, a mediocre dentist will probably be better than no dentist at all.

### 5.2.2 Fields of Expertise

An expert can only be relied on in areas that fall within her **field of expertise**. Someone who is an expert in one area needn't be an expert in other areas, and no one is an expert in every field. When a person gets outside her area of expertise, she often doesn't know any more than anybody else. For example, a good dentist is an expert on teeth, but probably not about the pancreas or the wishbone formation.

## 5.3 Evaluating Claims to Expertise

Appeals to experts can be thought of as arguments with the following form:

Premise: An expert on the subject says that  $X$  is true.  
Therefore,  $X$  is true.

Good sources:  
1. factual reliability  
2. personal reliability

Or, breaking the premise here up into smaller parts:

Premise 1:  $E$  is an expert on the subject  $S$   
Premise 2: Claim  $X$  involves the subject  $S$   
Premise 3:  $E$  says that  $X$  is true.  
Therefore,  $X$  is true.

We don't usually say this sort of thing explicitly when we cite an expert in an effort to convince others that  $X$  is true. But we usually go through this *sort* of reasoning (though usually not very consciously or explicitly) any time we rely on an expert.

Such arguments are not deductively valid. Why not? Because even the best experts can be wrong. But in the case of *legitimate* appeals to authority the argument will often be inductively strong. The better the authority, and the more the authorities agree about the issue in question, the stronger the argument will be.

When evaluating the claims or advice of a potential expert we should ask ourselves the following seven questions:

1. Do we care enough about the issue to try to evaluate those who claim to be experts about it?
2. Is the field one in which there even are experts?
3. Is the source an expert on the relevant issue?
4. Has the source been quoted accurately?
5. Is the issue one in which the experts are (mostly) in agreement?
6. Is the source's claim one that is very unusual or surprising?
7. Is there any reason to think that the source might be biased or mistaken in this particular case?

We will consider these matters in turn.

### Is the Issue Important to us?

What you don't know *can* hurt you

Information is often very valuable. What you don't know *can* hurt you. It can even kill you. Suppose tests show that you (or your ten month old daughter, or your sixty year old grandfather) have a serious form of cancer. Various treatments are available, each with its pluses and minuses. Here you would want to find out what the experts thought about the merits of various treatments before you made a decision. Indeed, it would be very sensible to get more than one opinion.

**Information costs:** costs in time and energy to obtain information

But information is not intrinsically valuable. More isn't always better. The trick is to know when you need information and when you don't (an even bigger trick is mustering the energy to go get it when you know you need it). Lots of information is useless. No normal person will feel the need to know the exact number of blades of grass in her front yard, at least not enough to count them. Lots of information isn't relevant to our concerns. Moreover, even when it is, there is usually a *cost* to acquiring it. It takes time and effort to read the relevant things and to talk to the relevant people, and in many cases the return on this investment is not large enough to justify the effort.

Information is also not intrinsically important. What information is important to you at a given time depends on your needs and interests at that time. If Wilbur makes a claim about OU's record in football ten years ago, you may doubt his recollection, but the topic probably doesn't matter enough for you to check. But *if* you have a lot of money riding on a bet about OU's record over the past decade, it would be important to find out if Wilbur really is reliable about such things. Whether this is so or not depends on *your* priorities and the particular situation.

It is possible to spend too much time gathering information that we don't need and too little time gathering information that we do need. But most of us err much more in the direction of obtaining *less* information than we need, rather than obtaining too much. So we'll focus on that here.

There are also cases where we need more information but there isn't time to get it. If you are driving down a country road and find someone who is bleeding badly as a result of a motorcycle accident, it would be very useful to know a good deal about first aid. But you don't have time to acquire the information you need. It is useful to acquire such skill in first aid, but there will always be situations we aren't prepared for, and here we just have to do the best that we can.

### Are there even Experts in the Relevant Field?

In some areas there may not be any experts at all. Are there people we can rely on to know whether capital punishment or physician-assisted suicide is sometimes acceptable or not? In extreme cases people claim to know things that couldn't



possibly be known, at least not now. No one, for example, can now know whether there is intelligent life on other planets (though some people are in a position to make more informed estimates about such things than others are). In cases where you can't be certain who, if anyone, is an expert, it is best to keep an open mind and remain undecided or, if you feel you must have some view, to accept an opinion very tentatively and provisionally.

### **Is the Source an Expert on the Relevant Issue?**

We can only rely on an expert on matters that fall within her **field of expertise**. So it is always important to ask whether the source is likely to be reliable about the subject matter at hand. Advertisements featuring celebrity endorsements often show people who aren't experts about the products they hawk. Michael Jordan probably knows a lot about nutrition, but he can't be relied on to know whether Wheaties are more nutritious than comparable cereals.

In most cases, though, the ad is probably not really designed to strike us as an appeal to expertise. It is so obvious that celebrities are rarely experts about the things they sell that we aren't likely to be taken in by them. Such advertisements are probably aimed more at getting us to identify emotionally with a product because we like (or want to be like) the celebrity selling it.

The views of famous people are often cited on matters where there are not experts. Albert Einstein is one of the greatest physicists who ever lived, and he did think carefully about many things. Even so, there is little reason to think that he discovered deep truths about religion or ethics. With the knowledge explosion, there are also people who *were* experts in a field, but who didn't keep up with things. If their information is sufficiently out of date, they may appear to be experts when they no longer are.

### **Is the Expert Quoted Accurately?**

It should go without saying that the source needs to be quoted accurately, but we often fail to do this. Usually no one bothers to check, and so a misquotation or inaccurate paraphrase easily escapes notice. Misquotation is sometimes intentional; it can be useful to cite some respected person to help make our case. But human memory is very fallible, and a misquotation often results from an honest mistake. It is also possible to quote someone accurately, but to take their remark out of context or to omit various qualifications they would make. Such quotations are also of little value in constructing a good appeal to expertise.

### **Are the Experts in Substantial Agreement?**

Experts in any area will disagree now and then, but in some areas disagreement is the typical state of affairs. For example, there is considerable disagreement among good economists on long-range economic forecasts or on the trends the stock market will take. There is little consensus among meteorologists about long-range weather forecasts. Able scientists disagree about the likelihood that there is life on other planets. When such disagreement is widespread, some of the experts are bound to be wrong, and we cannot reasonably expect that the expert we happen to rely on will be one of those who turns out to be right.

We can't expect total agreement among all of the experts, of course, so again we face a matter of degree (the more agreement among the experts the better). In cases where reasonable disagreement is inevitable, it is impossible to rely uncritically on experts, and you must obtain as much information as you can and think critically about the issues for yourself.

### **Is the Claim Unusually Surprising?**

When someone makes a claim that almost everyone agrees is true (e.g., that Norman is in Oklahoma), they don't need to build a case for it. Life is short, and we don't want to hear arguments to support everything anybody says. But if someone makes a surprising or controversial or implausible claim (e.g., that black U.N. helicopters are patrolling Yale's campus), then it is *their* responsibility to give reasons for their claim.

The more implausible the claim, the heavier their burden of proof. The basic point is that people can be wrong, and if a claim is extremely surprising, it may be far more *likely* that the source made a mistake (or is lying) than that the surprising claim is true (we will return to this topic when we discuss appeals to ignorance in the next module.)

### **Is the Expert Likely to be Biased or Mistaken?**

Experts are only human, and they are subject to the same biases and flaws as the rest of us.

#### **Vested Interests**

In some cases an expert may have a reason to deceive us. There is a sucker born every minute, and experts can use their credentials to take advantage of this. There will always be people with advanced degrees or years of training who offer a quick

fix in order to make a fast buck. Doctors working for tobacco companies did many studies that allegedly failed to establish the harmful effects of smoking.

If it is really obvious that someone stands to gain if we follow their advice, we are likely to be suspicious. But it isn't always clear when this is the case. For example, a skilled financial advisor often gets a cut if you invest your money in the mutual funds he recommends. The adjustor from the insurance company may well be an expert on roofs, but it can be in her interest to have the insurance company pay you as little as possible. Of course such people are often honest and do give good advice, but it is always important to know whether others have something to gain before we follow their advice.

### Honest Mistakes

Sometimes an expert has good intentions but is still prone to error for some reason or another. A referee who is usually good at telling whether a basketball player was guilty of charging may miss a call because she wasn't in a position to see the action clearly. A conscientious psychiatrist who is adept at spotting problems in adolescents may have a blind spot when it comes to her own children.

## 5.4 Who Do we Listen To?

We are more likely to believe someone we regard as an expert than someone we don't. In numerous studies people have been given a passage containing claims or arguments. Some of them are told that the passage is by someone they are likely to regard as an expert, e.g., a medical doctor writing in the *New England Journal of Medicine* or a professor doing biological research at Harvard University.

Other subjects are told that the very same passage was from a source they are not likely to regard as an expert, e.g., that it's a translation from *Pravda* or the latest offering on Wilbur's Home Page. People are much more likely to accept the claims and arguments when they are attributed to the more credible source.

There is nothing wrong with this. Better we should believe sources we don't find credible? But this very sensible phenomenon creates an opening for people who want to influence or manipulate us. We can't usually identify an expert solely on the basis of what she says—if we knew enough to do this, we probably wouldn't need an expert in the first place. We have no recourse but to rely on characteristics that often accompany expertise—characteristics like , title, the recommendations of others—that are good indicators of expertise. And so someone who appears to have these characteristics can pass themselves off as an expert even when they really aren't.

### 5.4.1 Faking Expertise: The Aura of Authority

#### Halo Effects

When a person seems to have one positive characteristic or trait, we often assume that they have other positive characteristics or traits. This is called the **halo effect**. One positive trait seems to set up a positive aura or halo of other positive characteristics around the person (we will study halo effects in Chapter 15).

It is legitimate to infer the presence of one positive feature on the basis of other positive features only if there is a strong, objective connection between the two traits (we will see later that such a connection is called a *correlation*. (p. 287). Some traits, like a person's title and or institutional affiliation, *do* correlate well with expertise. But the correlation isn't perfect, and sometimes people with an impressive title or prestigious job are not experts at all. Moreover some people will pretend to have characteristics that are good indicators of expertise in order to take advantage of us.

#### Titles

Medical doctors, lawyers, professors and so on have titles that often do signal expertise. Such titles also create such a strong halo that it extends to completely irrelevant things. In an experiment conducted in Australia a man was introduced to five different college classes as a visitor from Cambridge University, but different titles were attributed to him in the various classes. The titles were ones common in British and Australian universities. In the first class he was introduced as a student, in the second as a demonstrator, in the third a lecturer, in the fourth as a senior lecturer, and in the fifth as a professor.

When he left, the students were asked to estimate his height. With each step up the ladder of status, he gained about half an inch, so that when he was a professor he seemed to be two and a half inches taller than when he was a student.

Since titles can create a halo that involves quite irrelevant things like height, it is not surprising that they can create a halo that extends to expertise and, perhaps, even to things like honesty. But unfortunately, although years of training often lead to expertise they don't always lead to personal reliability. Late-night television is full of infomercials in which real doctors hawk quick fixes to help you lose weight, quite smoking, or get off the sauce.

It is also possible to fake having a title. Con men do it all the time. So when someone claiming an impressive title offers us advice, it is always worth asking ourselves whether they really do have the training they claim to have and, if so, whether there are reasons why they might be biased or mistaken in the current situation.

### Additional Indicators of Expertise

There are also **institutional halos**. We tend to think that members of prestigious institutions, e.g., Ivy League universities, are likely to be experts because of their affiliation. This is usually very reasonable, because institutional affiliation often is a sign of expertise. As with titles, however, it can be exploited by people with affiliations in an effort to make a fast buck, and it can even be faked by those who are skilled (and unscrupulous) enough to do so.

**Self-assurance and confidence** can also make a person's claims seem more credible. Lyndon Johnson used to say "Nothing convinces like conviction," and studies show that the more confident and certain a witness in a courtroom appears, the more believable others find her (though there is little correlation between confidence and accuracy). This seems to extend to experts in general. It is very easy, even natural, to think to ourselves: "That person wouldn't be sounding so sure of things unless they really knew, so ...". But as we will see later the correlation between confidence and accuracy is far from perfect.

**Clothes, jargon, non-verbal cues** (e.g., "body language") and other image-enhancing devices can also be used to create an aura of expertise. The clothes we wear serve as indicators of status, which we in turn use as an indicator of expertise. One high-status "uniform" in our society is the business suit. In one study it was found that people were over three times as likely to follow a jaywalker across a busy street if it was a man in a business suit.

It is no accident that people in commercials and ads often wear a white lab coat and sit in a book-lined study or an impressive looking lab. Such props create an atmosphere of expertise, and this can lead us to suspend careful and critical examination of their claims.

### Intimidation

A fake expert can sometimes do a snow job by using a lot of technical jargon. Indeed, even genuine, well meaning experts sometimes intimidate us with a barrage of technical terms. We are frequently reluctant to ask for an explanation that we can understand because we don't want to look ignorant or stupid. You may have experienced this on a visit to a doctor. She quickly describes what is wrong, often in words you don't understand, then confidently tells you what to do and hurries off to see the next patient.

In cases like this it is important to stick by your guns. There is no reason why we should understand the jargon it took experts years to master. You are the one paying the expert, so you have a right to hear their opinion in terms you can understand and to have them give you reasons for doing the things they recommend.

Some people find it easier to do this if they write out a series of questions in advance. And if it turns out that the expert isn't genuine (an unlikely event in the case of most doctors), having to explain their terms may expose their lack of knowledge.

### **Stereotypes**

Many people have stereotypes and prejudices which lead them to see members of certain groups as more likely to be experts than members of other groups. In our society, other things being equal, women are less likely to be seen as competent experts than men. For example, male and female groups are more likely to adopt the suggestion of a male than of a female. Stereotypes can lead people to see members of certain racial or ethnic groups, people with regional accents, and so on as less likely to have genuine expertise.

### **5.4.2 Appearing to Go Against Self-interest**

We all know that if someone stands to profit when we take their advice, we should think twice before taking it. But it is possible to exploit the fact that people without a vested interest are seen as more objective authorities. The trick is to appear not to be acting from self-interest while really doing just that.

Not long ago I was shopping for a new television. At my first stop the salesperson began by telling me that the Mitsubishi was the most expensive 27 inch television his store carried and that it was what the boss wanted him to push. Then, after looking around conspiratorially to make sure his boss wasn't within earshot, he confided that the Mitsubishi wasn't as good a bargain as a slightly cheaper model by Sony.

Although I don't know for certain what the salesman's motives were, the claim seemed well-calculated to show that he was honest and had my best interests at heart. After all, if he was simply out to make a fast buck, he wouldn't have clued me in about the defects of the Mitsubishi. But of course his commission wouldn't have been much less if I had purchased the Sony (which turned out to be the second most expensive model on the floor).

The psychologist Robert Cialdini took a job in a restaurant to study the techniques waiters used to maximize their tips. He found that the most successful waiter often used this sort of strategy. A large group would be seated. Then the waiter would tell the first person who ordered that the dish she asked for hadn't turned out very well that evening, and he would recommend something slightly cheaper. This would ingratiate him to his customers, sending the message that he was looking out for their interests, even at the cost of a larger tip to himself. But it turned out that this waiter's average tip was the highest in the restaurant.

A credible expert needs to have factual reliability *and* personal reliability. It is often possible to simulate the appearance of both, and many people make their living doing exactly that. In later chapters we will see how easy it is to do this, but if you think about it you can find plenty of examples in your own experience. The less we think about the things we hear, the easier it is to be a patsy. As always, the moral is to think; the more tuned in we are the easier it will be to think critically about what an alleged expert is telling us and the better we will be able to evaluate it.

## 5.5 Evaluating Testimony in General

We constantly rely on the claims of people who are not experts in any well-established field. We ask a stranger for directions to the Student Union. We ask friends and acquaintances which restaurants are worth going to and which are best avoided. If we are contemplating going out with someone, we might ask people who previously went out with the person how things worked out.

We also appeal to people in general; “They say that ...”. When we appeal to the fact that people in general (or people in some group we care about) think some claim is true, we are employing an appeal to authority called an **appeal to popularity**. Such appeals are common, and they can be very effective. In extreme cases, a **bandwagon effect** occurs when large numbers of people embrace a view (or support a cause like a particular political candidate) because other people have done so. Some people jumped on the band wagon, and since others don’t want to be left behind, they jump on too.

When we appeal to the fact that people have traditionally thought that some claim was true, we are employing an appeal to authority called an **appeal to tradition**. Such appeals may be legitimate, but they often are not. It depends on whether the group, either people now or people in the past, is a reliable judge about the issue in question.

Normally we believe much of what we hear, and unless there is a good reason not to that is entirely sensible. But anyone can be mistaken, and sometimes people lie. Furthermore, there are various pitfalls, including halo effects. Just as we are more likely to take the word of people who seem to be experts, we are more likely to change our views as a result of claims by someone we regard as similar to us. We will also see that one of the strongest halos is created by physical attractiveness, and it has been found that people are more persuaded by those they find physically attractive.

Most of the considerations that are relevant to evaluating the reliability of experts also apply when evaluating the reliability of your roommate or Aunt Sally

or even a stranger on the street. Indeed, we could regard such people as experts about some fairly limited subject matter like the restaurants in Norman or romantic interludes with Wilbur.

**The Seven Questions Revisited** We will quickly run back through the seven questions to ask about alleged experts and see how they apply to testimony in general, regardless of the source.

1. Do we care enough about this issue to try to evaluate the likelihood that a given source about it is accurate.

The general point here is the same whether the potential source is a world-class expert or just someone we meet on the street. But the costs of getting information from someone you know or encounter may be lower than the costs of getting information from an expert, so it may be reasonable to collect more information from those around you.

2. Is the issue one in which anyone can really be relied on to know the facts?

The point here is the same regardless of the source. If there are no experts in the field, then there is little likelihood that your friends and acquaintances will be particularly good sources of information about it.

3. Is the source generally right about this sort of issue?

Perhaps Anne has always given good advice on fixing computers, while Sam has often been wrong. Sally has always provided good advice about who to go out with, while Wilbur's advice is hopeless.

4. Is the issue one where people would mostly agree?

If there is little agreement among others about something, you are on your own.

5. Is the source's claim very unusual or surprising?

The point here is the same regardless of the source. If a claim is sufficiently unlikely it is more probable that the source is wrong than that the claim is true.



6. Is there any reason to think that the source might be biased or mistaken in this particular case?

Sally is a good judge of people and full of insights about their personalities, but she has a blind spot about Burt. Bill usually gives good advice, but he's really been stressed out lately. John saw the car I asked about, but it was dark and he could have made a mistake about its license number. Indeed, as we will see later in this module, even honest eyewitnesses are much less reliable than people commonly suppose.

7. Has the source been quoted accurately?

Hank tells us that Sally said that Cindy and Paul are back together again. Is there any reason to think Hank might be getting it wrong?

## 5.6 Safeguards

The following steps will help us spot, and so resist, fallacious appeals to expertise.

1. **Actively evaluate** claims and arguments that matter to you.
2. Check the alleged source's **credentials** and **track record**.
3. Check **multiple and independent sources** if the issue is important to you.
4. Determine whether it is in the experts's self interest to **deceive** us (e.g., is she trying to sell us something?).
5. Determine whether there is some special reason why she might be **mistaken** on this occasion.
6. Develop **your own expertise**.
7. Try to look at the issue from **multiple perspectives**.

**1. Tune In** One of the greatest obstacles to evaluating potential sources of information is that we often listen to them with our mind out of gear. Counterfeit authorities want us to follow their suggestions, and this works best if we go along, passively, mindlessly, without really thinking about what they are saying. Habit and routine and laziness encourage this. It requires an effort to think about things. But the more we do it, the easier it will become.

**2. Check the track record** Check the alleged expert's credentials and track record. If they have a history of making mostly true claims in a given area, that

gives us a reason to think they will be right in the future. Sometimes the track record is pretty clear. If several of your friends have had good experiences with a particular doctor when they had colds, it's sensible to go to her when you have a cold. If Tom's claims about which courses to take have always been wrong, he's not a good person to ask the next time around. The track records of many publications are reasonably good, whereas the track record of the *National Inquirer* is not. Checking a track record can be difficult, but if the issue is one that really matters to us, it is worth trying to do.

**3. Check multiple and independent sources** When you are uncertain whether an expert's claim is correct, it is prudent to check several sources. Get a second opinion (and, if the issue really matters to you, a third). But you must take care to find sources that are *independent* of each other. There is little point in checking several copies of today's *Washington Post* to be sure that the first copy was right. And if you ask six different people about the time of the final in Philosophy 101 but they all got their information from Wilbur (who misread the time in the syllabus), you will still be misinformed.

When *independent* sources agree, you can have more confidence in their joint testimony than you could in the testimony of any one of them alone. In many cases it is difficult to find multiple sources, but on the world wide web, where credentials and track record can be difficult to assess, finding multiple authorities is often quite easy. Indeed, you can use the net to get a second opinion after someone you know has given you a first.

**4. Consider Possible Biases** Ask yourself whether there is any reason why the alleged expert might be biased about this particular case. Do they have a financial stake in it? In the case of celebrity endorsements and infomercials the biases are usually obvious. The person stands to make a fast buck from us if we believe what they tell us. But in other cases vested interests may be less obvious. Indeed, in some cases the vested interest may simply be the desire to seem right.

**5. Consider Possible Sources of Error** Can you think of any reasons why the expert might make an honest mistake in this particular case? Might she have some sort of blind spot about it (as many of us do when it comes to our loved ones)? Were there reasons why the observations or tests might not be reliable (perhaps she is a good lab technician, but the police did a sloppy job gathering the DNA samples).

**6. Developing your own Expertise** In cases that really matter to you, you need, to some degree, to **become your own expert**. It is increasingly clear, for example, that people need to learn more about healthy life styles and how to manage their own medical conditions. In doing this we should of course rely on experts, but we have more first hand-knowledge about ourselves than others do, and we have a greater interest in obtaining accurate information about it.

**7. Look at things from Several Perspectives** One of the best ways to avoid flawed reasoning is to think about things from more than one point of view. This strategy is less relevant to testimony than it will be to some of the things we will study later, but it is still useful to try to put yourself in the position of the source you are evaluating. Can you think of other perspectives from which the experts claims would seem less plausible? Would you have any reason to make this claim if it weren't true?

In short, the key is to find *good authorities* who don't have any reasons to misrepresent the facts. If the matter is really important to us (as some medical questions are), we should also try to obtain several independent opinions from different sources. And we should always tune in when the topic is relevant to us.

As we work our way through the course we will find certain sorts of errors that we all tend to make. In cases where such errors are likely, it is important not to accept someone's claim too quickly. But before turning to errors, we will devote a chapter to one of the most important sources of information in today's world: the world wide web.<sup>1</sup>

---

<sup>1</sup>For statistics about the number of violent episodes the average child will see on television see D. Kunkel, et al., *The National Television Violence Study*, Mediascope, 1996. The study showing how estimates of heights were influenced by titles was conducted by W. F. Dukes and W. Bevan in "Accentuations and Response Variability of Personally Relevant Objects," *Journal of Personality* 20 (1952); 457–65. For a discussion of institutional halos see James S. Fairweather, "Reputational Quality of Academic Programs: The Institutional Halo," *Research in Higher Education* 28 (1988); 345–365. For evidence that confidence increases perceived credibility of witnesses see Philip Zimbardo and Michael Leippe's *The Psychology of Attitude Change and Social Influence*, Temple University Press, 1991, pp. 324ff. The jaywalking study is described in B. Mullin, C. Cooper, and J. Driskell, "Jaywalking as a Function of Model Behavior," *Personality and Social Psychology Bulletin*, 16 (1990); 320–330. For the study showing that groups were more likely to accept the suggestions of males see C. L. Ridgeway and C. K. Jacobson, "Sources of Status and Influence in All-female and mixed-sex groups," *Sociological Quarterly*, 18 (1997); 413–425. Robert B. Cialdini's observations of waiters, along with a detailed and fascinating discussion of influence in general, may be found in his book *Influence: The Psychology of Persuasion* Quill, 1993 (the example of the waiter is on p. 233). Evidence that we are more likely to change our views in response to claims made by those similar to us may be found in T. C. Brock, "Communicator-Recipient Similarity and Decision Change," *Journal of Personality and Social Psychology* 1 (1965); 650–654. Evidence that we are more likely to change our views in response to claims by those we regard as physically attractive may be found

## 5.7 Chapter Exercises

Answers to selected exercises are given on page 104.

1. For each of the following areas say
  1. Can there really be experts in this area? If not, why not?
  2. What sorts of people (if any) would be good experts in the field
  3. What sorts of people might bill themselves as experts about the topic, but not really be?
  4. How could you try to determine whether an alleged expert in the area is a genuine expert (and, if so, how good an expert they are)?
1. The way TV sets work
2. College football
3. College football recruiting
4. Shakespeare's plays
5. Who's two-timing who in Belleville, Kansas [population 2200]
6. The issue of whether gay marriages should be legalized
7. The effects of your astrological sign on your behavior
8. The Existence of God
9. Whether or not creationist theories about the origin of the universe are true.
10. How soon al-Qaeda will launch another major attack.
11. The artistic value of 1950's rock-n-roll
12. The precise number of people who lived on earth exactly 100,000 years ago
13. The issue of whether an accused murderer was really criminally insane at the time she allegedly took the pickaxe to her victims
14. The safety of nuclear power plants
15. The morality of abortion
16. Losing weight and keeping it off
17. Whether gun control is a good thing or not
18. Which majors are most likely to get well paying jobs when they graduate from college.

The following passages contain appeals to authority. Say whether the appeal is legitimate or not, and defend your answer.

---

in S. Chaiken "Communicator Physical Attractiveness and Persuasion," *Journal of Personality and Social Psychology* 37 (1979); 1387-1397.

1. Most people who teach critical reasoning are very skeptical of astrological predictions. But people have been using the stars to make predictions for hundreds and hundreds of years. They surely wouldn't do this if there weren't something to it.
2. Recent polls show that a large number of people believe in the power of astrology and the accuracy of astrological predictions.
3. According to Einstein, the idea of absolute motion is incoherent. And that's good enough for me. (This is one where background knowledge is needed.)
4. Both the Surgeon General and the American Heart Association insist that smoking is a leading cause of heart attacks. So it's a good idea to quit smoking.
5. The following passage appeared in Phil Dalton's column in *The Oklahoma Daily* (9/24/97, p. 4); how plausible is it?

I am not arguing for the legalization of marijuana. Instead, hemp should be legalized to help protect forests and woodlands and our rivers. (And yes, I did get all this information from the National Organization for Reform of Marijuana Law Website).

6. Suppose that U.S. Senator John McCain, who champions campaign finance reform, is giving arguing that the huge campaign contributions that large companies give to political candidates lead to a substantial amount of corruption in the American political system (e.g., by influencing which laws get made). What questions should you ask in order to evaluate his argument? How might you go about finding information supporting the other side (that it really doesn't lead to much corruption)? How would you evaluate what you hear or read on this issue?
7. Give an example (of your own) of an area where there are experts, but where it is likely that they will frequently disagree with one another. If you had to make a decision that required you to know something about the field, what would you do?
8. Give an example of your own of an area where there really do *not* seem to be experts at all. If you had to make a decision that required you to know something about the field, what would you do?
9. The average child will have seen at least 8000 murders and 100,000 other acts of violence depicted on television before they graduate from elementary school. Suppose someone uses this to argue that we should restrict violence on television. What sorts of information would you need to evaluate their argument? Could you

get it without relying on others? What people would be likely to have accurate information about the matter?

10. Wilbur and Wilma are discussing capital punishment.

1. Wilbur argues that we should have capital punishment because it deters terrible crimes like murder (i.e., it tends to keep people from committing murder). How might we decide whether or not she is right? Are there any experts who might have useful information on the matter? If so, what sorts of people are likely to be experts here?
2. Wilma counters that we should get rid of capital punishment because it is morally wrong. How might we decide whether or not she is right? Are there any experts who might have useful information on the matter? If so, what sorts of people are likely to be experts here?

11. Give an example of a celebrity endorsement. Do you think that such endorsements are an effective way of advertising? If you think that they are, explain why you think they work.

12. Suppose that you wanted to know about the long-term behavior of the stock market, so that you could begin investing a modest amount of money now, while you are still a student. Are there people who would know more about this than you do? If so, who? Are these people likely to be experts. If not, why not? If so, how might you try to check the claims of one of the experts about how you should invest your money?

13. It's pretty clear that the photographers caused Princess Di's death. I saw a long account of it two of the supermarket tabloids at the checkout counter at Homeland. Buy that?

### **Answers to Selected Chapter Exercises**

In many of the following cases there is no one right answer, but some answers are certainly better than others (and many possible answers are wrong).

1. The way TV sets work
  - (a) Likely experts: people who repair TVs, scientists and engineers who design TVs, some (though not all) people who sell TVs.
  - (b) Some salespersons act like they have more expertise than they do.

- (c) Salespeople do often have something to gain by getting you to buy a TV. Other things being equal, you would probably trust a salesperson who doesn't work on commission. But there is a great deal of variation among salespeople, and you cannot make a blanket generalization about them.
2. College football
- (a) Likely experts: College football coaches, sportswriters and sportscasters. Even the experts will disagree about some things here, but coaches who are successful year in and year out have something going for them.
- (b) Monday morning quarterbacks consider themselves experts.
3. College football recruiting
- (a) Sports journalists, high school coaches, and perceptive fans may have a good idea about recruiting, but the best experts here are probably good recruiters.
- (b) Fans who consider themselves experts.
4. Shakespeare's plays
5. Who's two-timing who in Belleville, Kansas [population 2200]
- (a) The point of this example is to emphasize that there are experts on all sorts of things. If you grew up in a small town (I *did* grow up in Belleville), you will remember town gossips who had all the dirt on everybody.
- (b) Some gossips like to pretend they have the dirt even when they don't. If you want accurate rumors, it's best to find a reliable gossip (of course as tabloid journalism attests, rumors are often more fun when they are inaccurate).
6. The issue of whether gay marriages should be legalized
7. The effects of your astrological sign on your behavior

When we get to a later module I'll try to convince you that there are no experts on this, because the planets and stars has no discernible impact on your character or behavior. Astrology is a pseudoscience.





## Chapter 6

# The Net: Finding and Evaluating Information on the Web

**Overview:** The internet is a new and unparalleled source of information, but it can be difficult to track down what want because there is so much information on the net and there is so little quality control. In this chapter we learn how to search efficiently for information on the web and how to evaluate it once we find it.

### Contents

---

<b>6.1</b>	<b>The World Wide Web</b> . . . . .	<b>108</b>
6.1.1	What is the World Wide Web? . . . . .	108
6.1.2	Bookmarks . . . . .	109
<b>6.2</b>	<b>Search Engines</b> . . . . .	<b>110</b>
6.2.1	Specific Search Engines . . . . .	111
6.2.2	Metasearch Engines . . . . .	112
6.2.3	Specialty Search Engines . . . . .	113
6.2.4	Rankings of Results . . . . .	113
<b>6.3</b>	<b>Refining your Search</b> . . . . .	<b>113</b>
<b>6.4</b>	<b>Evaluating Material on the Net</b> . . . . .	<b>115</b>
6.4.1	Stealth Advocacy . . . . .	119
<b>6.5</b>	<b>Evaluation Checklist</b> . . . . .	<b>120</b>
<b>6.6</b>	<b>Citing Information from the Net</b> . . . . .	<b>121</b>
<b>6.7</b>	<b>Chapter Exercises</b> . . . . .	<b>122</b>

---

## 6.1 The World Wide Web

Problems with the Net:  
1. Information Overload  
2. Quality Control

The internet is a new and unparalleled source of information. With a computer and a connection to the net, you can find information about almost anything. But there are two problems you face as soon as you start looking.

1. There are now several hundred million pages on-line, and about a million new pages are added every day. There is so much information that locating what you need can be like looking for a needle in a haystack.
2. Anyone with a computer and a connection to the net can set up their own web site, so it is not surprising that there is almost no quality control. Even once you find information on the topic you want, it can be difficult to judge how accurate and complete it is.

In this chapter we will learn about the net and ways to solve these two problems.

### 6.1.1 What is the World Wide Web?

The *internet* is worldwide network of computers. Once you connect to it, you can send and receive information practically instantly. The initials, WWW (or W3), stand for the *world wide web*. The web is the constantly changing collection of web documents (sites) that you can access once you are on the internet.

### HTML

Documents on the web are written in a “marked-up” version on English called HTML (hypertext markup language). An HTML document is just a plain-text (also known as ASCII) file that contains codes or tags that determine how a web page will look when it is displayed on a piece of software called a browser. For example, to italicize a word, you enclose it in the tags for italics, `<I>like so</I>`.

But don't worry. You do not need to know anything about these tags in order to use the web, because a browser, which is just a program for displaying HTML on screen, will handle that for you. The most popular browsers are Netscape and Internet Explorer. In addition to displaying text they allow you to access newsgroups, display images, play sound files, and more.

### Hypertext

The HT in HTML stand for **hypertext**. Hypertext is text that includes **links**. Links are words or phrases, which will usually be underlined or displayed in a different color from the rest of the text, that will take you to another spot in the document

you are viewing or to another document entirely—even one at some site half way around the world. The main reason HTML is so useful is that it allows you to construct a document that is linked to text or images in other documents. When you encounter a link, you can click on it and jump to the linked document.

Figure 6.1 shows how the home page for a recent section of my critical reasoning course looks in the Internet Explorer 6.0 browser (it looks very similar in Netscape or in other browsers like Opera). The underlined words and phrases are links. You just move the mouse until the arrow or pointer is on top of a link you want to follow and then click the left mouse button. Links are typically a different color from the rest of the text, so they are usually easy to locate.




Figure 6.1: Viewing a Page in a Browser

### 6.1.2 Bookmarks

If you visit a web site frequently it becomes tiresome to type out its complete address every time. To avoid this, browsers let you set **bookmarks** that let you return to frequently visited sites with a click of a mouse. Your bookmark file is like an electronic address book. In Netscape there is a button at the top labeled *Bookmarks*. In Internet Explorer there is a button with the same function, but it is labeled *Favorites*. When you are at a site that you know you'll want to revisit,

**Bookmarks:** let you return to frequently visited sites with a click of the mouse

simply click on the *Bookmarks* or *Favorites* label. The browser will ask you if you want to add the address for this site to your bookmark file. Just click *Yes*. When you want to revisit the site, click on the bookmarks label and the address will be there. Move the cursor to this address, click the mouse, and you will be returned to this site.

## 6.2 Search Engines

**Search engine:** a program for finding information on the web

It can be fun to surf the net, wandering from one site to another, going wherever the links take you. But if you are trying to find information about a specific topic, e.g., what courses are offered at OU next semester, the starting salaries of people who graduate with various majors, or OU's record in bowl games, you need some way to sift through the endless information on the net. You need a search engine. A search engine is a piece of software that helps you find things.

A search engine does not go out and search the web each time you type in a word or phrase. It is a service that indexes or stores a huge amount of information about the contents of many web sites. This information is stored in a "database." The database of a search engine contains a list of all the words in all the web pages that the engine knows about. If you type in a keyword or phrase, e.g., Bill Clinton, the search engine will consult its database and give you a list of links to sites that contain information about Clinton.

A search engine has a database, a record of information about various kinds, that it consults when you send it an inquiry. For example, it will have a record of all of the web sites that it knows about which contain the phrase 'Critical reasoning'. Different search engines obtain their databases in different ways. Some consist of indices built by people who search the web looking for information about various topics. Others use software to go out and see what is on the net. But you don't need to know about how any of this works.

Different search engines have different databases, and so different engines are better for different purposes. When a search fails to turn up information that you are searching for, try another search engine. All you need to do to use a search engine is to type in its URL in your browser, and it will handle things from there. The URL (uniform resource locator) is just the web address. For example, the URL for this course's page is <http://www.ou.edu/ouphil/faculty/chris/critreas.html>.

Figure 6.2 on the next page shows the AltaVista Search Engine (using the Internet Explorer 5.0 browser; it would look very similar with Netscape). The white arrow points to the spot where you type in the words or phrases that you want to search for. The phrase 'Critical Reasoning' and 'OU' have

been typed in this example (you will learn what this phrase means below). When you get to this point, you just hit the Enter key or click your mouse on the Search button to initiate your search.



Figure 6.2: Searching with AltaVista

There are a number of good search engines, and new ones are introduced frequently. Here are a few of the best. Don't be intimidated by the length of this list. All search engines work in pretty much the same way; once you've seen one, you've pretty much seen them all.

### 6.2.1 Specific Search Engines

#### **AltaVista** <http://www.altavista.com>

This is the largest search engine. You type in a keyword or phrase, and it displays all the sites it knows about that pertain to that word. It will include a brief description of the document at that site, the date it was last modified (this is important if you need current information), and a link you can click on that will take you to that site.

**Google <http://www.google.com>**

Google employs a more sophisticated relevance ranking of sites than most other search engines (it's based on the number of links in related documents). It returns a fairly high ratio of useful information to junk, and is good for academic material

**Yahoo <http://www.yahoo.com>**

Yahoo organizes its sites into fourteen main categories (e.g., Arts, Humanities, Sports). If you aren't sure how to begin, Yahoo is organization by topics may help you get started. It's good, but not as comprehensive as AltaVista.

**HotBot <http://www.hotbot.com>**

Easy to use and has very strong search capabilities

**Infoseek <http://www.infoseek.com>**

Infoseek has over 50 million web pages in its database and is very accurate.

If you can't find something with one search engine, try another. But most of them will work well for common topics, and there is something to be said for focusing on one or two engines so that you get used to them.

You can also go to various web sites which contain links to many search engines on the same page. Then just click on one of them to fire up the search engine you want. You can find such lists in a number of places; here are a three typical ones:

1. <http://www.home.co.il/search.html>
2. <http://www.phoenixgate.com/search.html>
3. <http://search.tht.net/>

### 6.2.2 Metasearch Engines

A **metasearch engine** will take your entry and actually use several different search engines to find hits for your keyword. Good metasearch engines include:

**AskJeeves <http://www.ask.com>**

This lets you ask a question in plain English, such as "Where can I see tax returns from President Bill Clinton?" or "Where can I find a map of Norman, Oklahoma?"

**Dogpile <http://www.dogpile.com>**

Dogpile employs all of the major search engines.

**6.2.3 Specialty Search Engines**

There is an increasing number of special search engines for finding information about specific fields. You can find a list of many of them at the below sites (the first one is especially good):

- **Speciality Search Engines:** <http://www.leidenuniv.nl/ub/biv/specials.htm>
- **Search Engine Watch:** <http://www.searchenginewatch.com>
- **dejanews.com**, a search engine for newsgroups: <http://www.deja.com>

You can find very accessible information on search engines, including on-line tutorials, at

- <http://www.library.fullerton.edu/interws.htm>

**6.2.4 Rankings of Results**

Most search engines return a list of sites ranked in order of relevance to the words that you used in your search. The first listing will be the most relevant, the second listing the second most relevant, and so on. But the techniques that search engines use for this are highly fallible. For example, some search engines rank a page higher the more times it contains the word that you searched for. This will sometimes help you zero in on what you want, but you can't count on it. You will have better luck finding things if you refine your search.

**6.3 Refining your Search**

There is so much information on the web that it is rarely useful to search for a single commonly used word or phrase. If you type in Bill Clinton, you will get millions of "hits." So you will need to refine your search. You can do this by asking for sites that refer to more than one topic (e.g., Bill Clinton and Monica Lewinsky). You can also exclude irrelevant topics.

Different browsers use different conventions for refining searches, but the basic ideas behind the different conventions are similar. You can try a search for a keyword, and if that returns too many hits, you can refine your search. This takes a little practice, but once you get the hang of it, it's not difficult. Searching the net is not an exact science, but with fifteen or twenty minutes practice, you'll become fairly proficient.

### Narrowing Your Search

There are millions of sites out there, and the biggest problem is rarely that you can't find any sites that discuss things you are interested in. The problem is usually that you find so many sites that you are overwhelmed by all the information. The best way to deal with this problem is to be as specific as you can. If you want to find out about Doberman dogs, type `Doberman` rather than `Dog`.

If you are interested in specific information (e.g., how many phone calls did Linda Tripp tape), then you need to make your request specific. If you simply type `Linda Tripp`, you will get links to many sites, but lots of them won't contain the information that you want. You will often have to use a little trial and error to word your search in the best way.

If your target involves a sequence of words, put the sequence in quotation marks. This will lead most search engines to display just those sites that contain the phrase. For example, if you want to find out about Bruce Lee, put the entire phrase in quotes: `"Bruce Lee"`. If you don't, the search will return all the sites it knows about that refer to any old *Bruce* and all those that refer to any old *Lee*. If you want to find discussions of a book or movie, put the entire title in quotes, e.g., `"I know what you did last summer"`.

Many search engines will also let you put `and` and `near` between items in quotes. Thus `"Bill Clinton" near "affairs"` will bring up all the sites that discuss Bill Clinton and affairs (this will still be a very large number of sites). If you want to find recipes for chocolate chip cookies try `"chocolate chip cookies" AND "recipe"`. In many search engines you can get the same effect by putting a plus sign, `+`, in front of each word, as in `+Clinton +affair`.

Many engines have additional features, and some work in slightly different ways, but these pointers will be enough to get you started. All good search engines also have on-line help, so if things aren't clear, click on the engine's `HELP` button. One of the best ways to learn about computers, and this is especially true about using the net, is to approach them in a playful, inquisitive spirit and just try things out. Computers give you immediate feedback, so you can learn quite quickly how to find things on the net.

### Boolean Operators

**AND**, **OR** and **NOT** are known as Boolean Operators (for historical reasons that won't matter here). All decent search engines let you search for an exact phrase you typed, all the words in the phrase but not necessarily together, any of the words in a phrase, etc. In this subsection we will see how to use Boolean operators to do



this (you already know about the first one, AND).

**AND** As noted above, to search for two (or more) words on the same web page type the word **AND** between them. AND means just what you would expect; it corresponds to a conjunction in logic and intersection in set theory. For example: Clinton AND scandal (in many browsers the + sign does the same thing: +Clinton +scandal).

**OR** To search for either (or both) of two words on the sample page type an OR between them. OR corresponds to a disjunction in logic and to union in set theory. For example: doberman OR collie will return all pages that mention either dobermans or collies, or both.

**NOT** To search for pages that include the first word but not the second, type NOT before the second word. Not corresponds to negation in logic and to complements in set theory. For example, if you are interested in facts about the sabbath you probably won't want to visit sites dedicated to the rock group Black Sabbath. To avoid this, type Sabbath AND NOT Black. In many browsers you can achieve the same effect with a minus sign: Sabbath -Black.

To search for various forms of a word, type an asterisk at the end. For example, cat\* would return pages with the words *cat*, *cats*, *catty*, etc. You can also string various items together, e.g., ``Bill Clinton`` AND ``scandal`` AND NOT ``Whitewater`` would find sites that discuss Bill Clinton and scandals other than Whitewater.

## 6.4 Evaluating Material on the Net

Anyone with a few dollars and access to a computer can set up a web site. So it is not surprising that a lot of the information on the net is inaccurate. Some of the material is merely the author's opinion, often a quirky opinion, some of it is plagiarism, some is just a mishmash cobbled together from other sites that are never cited. Other information may have been accurate when it was posted, but many websites are not updated and they may support their claims with links to sites that no longer exist (this is known as link-rot).

So how can you tell whether a given site is reliable? In evaluating any would-be authority, it a good idea to run through seven questions that are always relevant to evaluating the claims of other people, the media, etc.. But there are special things to look for when asking these questions about claims that you find on the web.

**Who is the Author?**

The first thing to ask is: who is the author? It isn't always easy to answer this question. Sometimes you find a page with no name on it. Sometimes you find a page with a name, but it's not clear that the person who posted the page actually *wrote* the material (there is a lot of borrowing and some plagiarism on the net). If you can't determine who the author is, it doesn't make sense to put much trust in the material.

**Who Maintains the Site?**

It is also important to ask who maintains the site. The *person* who oversees the details of the site is known as the webmaster, but often sites are maintained by organizations, e.g., businesses, university departments, advocacy groups, government agencies, and the like. Knowing who the site belongs to can alert you to the quality of information posted there and possible biases. You can often learn a good deal (for sites in the U.S.) by looking at the URL, the page's web address.

1. If the URL ends in `.edu` it is an educational site.
2. If it ends in `.com` it is a commercial organization (a company or business).
3. If it ends in `.org` it belongs to a nonprofit organization.
4. If it ends in `.gov` the site belongs to some governmental agency.

All these groups can have their special biases. Businesses are in the business of selling services and products and many organizations are in the business of selling ideas. But governmental agencies and educational institutions can also have agendas.

**Is the Author a Reliable Authority on the Topic?**

Once you discover who maintains the site and who the author is, the next question is whether the author should be considered a reliable source of information about the topic you are interested in. If the writing is sloppy, full of bad grammar and bad spelling, the author is probably an amateur. But the most basic question here is whether the author has any special training or experience that gives us any reason to think that he is a reliable source of information on the subject he's writing about.

Titles and professional affiliations are not infallible indicators of accuracy, but when you can find so much information on so many topics, they are a very important guide about quality of information. An article on heart disease written by a medical doctor and posted at a medical school's web site is much more likely to be accurate than an article posted by someone who provides no evidence that he knows what he is talking about.

In cases where you can't discover the credentials or track record of the author, ask whether there are any particular reasons to think that he or she would be reliable on this issue. What reasons or data or evidence does the author give to support her claims?

If the author provides documentation or references to support her views, e.g., by citing articles or books or by providing links to other sites, that gives us some reason to put confidence in her claims. But of course the same questions arise about the books or web sites the author cites, so citations alone aren't enough (if Wilbur cites his friend Sara's page and they worked on their pages together, her page shouldn't make Wilbur's page seem more reliable).

### Consult Independent Sources

One of the best ways to check on the accuracy of source is to consult additional sources (get a second opinion). With books or television, this can take too much time. But on the internet it is easy; it usually just involves clicking on several of the items your search engine pulls up. So if the information is important to you, consult *several sources*. And, as always, look for *independent sources*. If Wilbur's page just repeats the information on Sara's page, then his page doesn't provide independent backing for her claims.

### Biases?

As always, we should also ask whether the source is likely to be biased or mistaken about the topics discussed on her page. If they are selling something, we'll want to take their claims with a grain of salt. Often this is obvious, but not always. Just remember that if you are reading about a product at a site maintained by the people who produce it, you are usually reading an advertisement (it may be informative, but it's there to get you to buy it).

### Stealth Advertisements and Skewed Searches

A special problem on the internet is that the line between advertisements and the presentation of information (which is relatively clear cut in many newspapers and on the major television networks) is getting more and more blurred. On February 26, 1999 the *New York Times* reported that if you do a search to find a good Mexican restaurant in Los Angeles, for example, you are likely to turn up the *Los Angeles Times* listing of restaurants. It sounds like a neutral and objective source, but in fact it gives favorable placement of ads to advertisers.

A similar problem recently surfaced with Amazon.com, a huge internet book store, prominently displays a list of recommended books on their home page. It

turns out that the publishers of books listed there often paid Amazon.com (these are known as “co-op placements”). After complaints from customers, Amazon.com agreed to provide information about which book promotions were subsidized by publishers (although this information is not in a place where readers are likely to notice it). This sort of situation is not atypical. You know when you visit the site of a company or business they will be plugging their product, but they do so up front. Unfortunately, many other sites that look neutral are also plugging products under the guise of providing information (about restaurants, books, and the like).

As we noted above, many search engines have techniques to rank the results of your search in order of relevance to your original query. But some of these techniques are easily exploited. For example, some engines rank a page higher the more times it contains the word that you searched for. So people who want you to visit their site can ensure that it will come out high in a ranking by using a commonly-searched-for word over and over in their page. For example, if a car dealer’s site started off with the words “Cars, cars, cars, cars” it would come up high in the rankings of the sites that would be listed if you began your search with the word `cars`.

### Advocacy Pages

Authors and webmasters may be biased in other ways. It is very easy to find sites maintained by true believers in all sorts of causes, political movements, and the like. The pages at such sites are sometimes called **advocacy pages**. These are web pages that aim to influence opinion or promote causes; they are selling a product just as surely as businesses are, but the product is an idea or a political candidate or a cause.

There is nothing wrong with advocating a view or a cause, and if an author is forthright about what they are advocating, you can evaluate their claims in light of that knowledge. In many cases an advocacy page will provide accurate information about the nature and aims of the group that maintains the page. But it is naive to accept their factual claims, particularly their claims about opposing groups, uncritically. At the very least, you should check more neutral sites to verify their claims.

The web sites for major political parties and political candidates are good examples of advocacy pages. For a somewhat different example, consider the following passage, which appeared in Phil Dalton’s column in *The Oklahoma Daily* (9/24/97, p. 4)

I am not arguing for the legalization of marijuana. Instead, hemp should be legalized to help protect forests and woodlands and our

rivers. (And yes, I did get all this information from the National Organization for Reform of Marijuana Law Website).

These claims about hemp may very well be true, but if you cared enough to find out, you would want to verify them using more neutral websites.

### 6.4.1 Stealth Advocacy

Larger problems arise with **stealth advocacy pages**, pages at sites that promote a cause or viewpoint without acknowledging that this is their aim. If their authors have gone over the deep end it will often be obvious from the extreme nature of their claims (“All of the proponents of the other side are clearly deluded because . . .,” “It’s obvious to anybody in their right mind that . . .”). But this isn’t always the case, so it’s important to know whether the source has a particular axe to grind. When you suspect that this might be the case,

**Stealth advocacy:**  
promoting a view while  
appearing to be neutrally  
presenting facts

1. Read the page carefully to see if it is written from a particular “point of view.”
2. Ask yourself why the page was posted in the first place. Was the point simply to inform people, or does there seem to be an ulterior motive.

Information can be biased in various ways. Sometimes it is just blatantly false. That may not be too hard to spot. But it is also to be misleading by only telling part of the truth. A site might correctly but selectively quote statistics that favor their point of view while omitting equally correct statistics that support the other side. For example, a page claiming that O. J. Simpson didn’t kill his wife might correctly cite the statistic that only a very small proportion of husbands that beat their wives kill their wives (this is true, because a very low percentage of people are murdered). But the proportion of husbands who kill their wives is much higher among abusive husbands whose wives are later murdered.

It is also possible to present statistics or other data that are accurate, but to *interpret* them in questionable ways. You won’t usually know enough to evaluate detailed statistics, so it’s sensible to be mildly skeptical of information supplied by a site that is likely to be biased. The best policy here is to look for a neutral site or at least to look at a site maintained by the other side.

A special difficulty is that if you are already predisposed to think that a given answer to some controversial view is true, then you can always find sources on the net that seem to support your view. But finding a bunch of biased sources to support a biased view doesn’t make that view unbiased.

### **Currency**

Finally, if you need very current information on a topic, check to see when information was posted at a site. It may have been current when it was first put on the net, but many sites are rarely, or never, updated.

### **Ease of Use**

Even a reliable, well-documented site may be difficult to use because it is badly organized, too superficial (or too detailed), pitched at the wrong level for you, or is just plain confusing. Different styles of presentation will be appropriate for different readers. If a physician posts a technical discussion for other physicians most of us would be unable to follow it. This doesn't mean there is anything wrong with the page—it may be optimally written for its intended audience—but it won't be at the right level for most of us.

People often underestimate how much the style of presentation matters. You might read a page by one author that leaves you badly confused but learn a lot from another presentation that covers the same material in a clearer way and at a level that is more suited to your current knowledge.

It will be easier to remember these points if we summarize them in a checklist.

## **6.5 Evaluation Checklist**

### **1. Author**

- (a) Did person who posted page write its contents? If not, who did?
- (b) What reasons to think author is a qualified, reliable source?
  - i. Occupation, degree, credentials, experience, cited by trustworthy source?
  - ii. Can you find biographical information on the author at an independent site?

### **2. Site Owner**

- (a) Is the site maintained by a reputable organization (a University Department, well-known business, etc) or by an individual?
- (b) Is the site owner likely to be reliable about information of this type?
- (c) Any reasons to think site maintainer biased or has an agenda?

### **3. Content**

- (a) What evidence (if any) is presented for the claims?

- (b) Does the author explain *how* he gathered his information (e.g., surveys vs. gut feelings)?
- (c) Are references (bibliography, links to other sites) given so the claims can be checked?
- (d) Any reasons to think author biased?
  - i. Point of site: is information posted to inform or to persuade?

#### 4. Independent Sources

- (a) Does the information at independent sites agree with this?

#### 5. Currency

- (a) When was page first posted?
- (b) When was it last updated (you can view the directory the page is in to determine this)
- (c) When was the information I care about last updated (often hard to tell)

#### 6. Ease of Use

- (a) Is the information clear?
- (b) Is it sufficiently detailed and comprehensive?
- (c) Is it written at the right level for me?

## 6.6 Citing Information from the Net

There are no established conventions for citing information from the net, and different people or groups may prefer different formats. We won't be interested in questions about format here, but will instead think about the *point* of citations in order to see what *sorts of information* need to be included. The point of a citation is to provide objective information about your source so that other people can check it out. This can be difficult when the source is a web page, because web sites come and go; by the time you mention a site to someone else, it could be gone. So you need to indicate when you visited a cite. You should cite as many of the following as you can (if you can't discover most of these things, that in itself is a good reason to be leery of the site).

**Author** Who wrote the things on the site?

**Maintainer** Who maintains the site (in the case of a personal web page, this may be the same as an author, but it may also be a business, political group, charity organization or the like)?

**Credentials** What reasons do you have to think author is a reliable source of information?

**URL** Give the *full* web address of the site

**Original Posting** When was page first put on the web?

**Last Update** When was the site last updated (this is important, but fallible information, since even if parts of the page may have been updated the information you cite may not have been).

**Date Site Consulted** When did you visit the site?

In cases where you really need to rely on information from the site, email the webmaster if you have questions and *print the pages* (including references) containing information that you need. Information from a page that no longer exists is often better than no information at all, but you can't expect others to find it as convincing as a source that they can check for themselves.

## 6.7 Chapter Exercises

The only way to learn how to search the web is to do it. This exercise requires you to go on an internet scavenger hunt. Use at least three different search engines or metasearch engines listed above. Before you begin, think about which words or phrases to use in your search. Remember to be as specific as possible. If your first attempt is not successful, refine your search using the tips above. And if you can't find the answer with one search engine, try another. Write down the word or phrases that returned a hit, list the search engine you used, give the URL (web address) of the page containing the requested information, and attach a print out of the first page of the site.

### Internet Scavenger Hunt

1. Find a picture of Dick Chaney.

Search Word(s): \_\_\_\_\_

Search Engine: \_\_\_\_\_

URL: \_\_\_\_\_

2. Who is the highest paid player in the National Football League? How much did he make last year?

Search Word(s): \_\_\_\_\_

Search Engine: \_\_\_\_\_

URL: \_\_\_\_\_

3. Find a Dilbert Cartoon.



Search Word(s): \_\_\_\_\_  
Search Engine: \_\_\_\_\_  
URL: \_\_\_\_\_

4. Find a recipe for chocolate chip cookies.

Search Word(s): \_\_\_\_\_  
Search Engine: \_\_\_\_\_  
URL: \_\_\_\_\_

5. Where is Kosovo? Find a current map that shows its location relative to Albania and Macedonia. Is it a country or a region?

Search Word(s): \_\_\_\_\_  
Search Engine: \_\_\_\_\_  
URL: \_\_\_\_\_

6. Find a street map that shows where you live during the school year.

Search Word(s): \_\_\_\_\_  
Search Engine: \_\_\_\_\_  
URL: \_\_\_\_\_

7. Find a listing of the different human blood types and a brief description of each.

Search Word(s): \_\_\_\_\_  
Search Engine: \_\_\_\_\_  
URL: \_\_\_\_\_

8. Many cities have experienced recent drops in crimes. Has there been a drop in the rate of any major crimes in the city or town where you grew up? If so, which crimes?

Search Word(s): \_\_\_\_\_  
Search Engine: \_\_\_\_\_  
URL: \_\_\_\_\_

9. Where was Osama bin Laden born and where did he receive his schooling?

Search Word(s): \_\_\_\_\_  
Search Engine: \_\_\_\_\_  
URL: \_\_\_\_\_

10. What are the symptoms of breast cancer (or prostate cancer)? What, if anything, can people do to avoid these types of cancer?

Search Word(s): \_\_\_\_\_

Search Engine: \_\_\_\_\_

URL: \_\_\_\_\_

11. Use the web version of your library's homepage to get the information about the book *The Golden Bough* (who wrote it, when, what is the call number?).

12. How much money did Bush and Gore spend in their race of the Presidency?

13. Use the information from your internet scavenger hunt to evaluate the sites you found with respect to each of the items on the evaluation checklist above.

- Who is the highest paid player in the National Football League? How much did he make last year?
- Find a description of the different human blood types.
- Has there been a drop in the rate of any major crimes in the city or town where you grew up? If so, which crimes?

14. How many people in the United States are attacked by sharks each year? How many are killed by sharks?

15. What sort of business was the company Enron in? How much money did its stockholders lose?

16. What does the name 'al-Qaeda' mean? What language is it? How did the terrorist alliance get this name?

17. Find (and print) two pages that deal with the same topic. Evaluate the relative merits of the two sites, particularly their reliability.

18. Find (and print) a page that provides information that you think is very *unreliable*. Explain why you think it may be inaccurate.

19. Use a search engine to find three sites that discuss the problems of assessing accuracy of internet pages (do these pages themselves seem reliable)?

## Chapter 7

# Memory and Reasoning

**Overview:** Our memories do not store exact copies of things in the way that video tapes or CDs do. Memories are not fixed, inert encodings of information. Memory is *active*. It fills in details in an effort to make sense of things; it involves elaboration and reconstruction. This filling in is akin to inductive inference. The way we fill in the gaps, and even rewrite the past, is influenced by our expectations, emotions, and other features of the context in which we remember something. Memory plays a key role in all our thought, and so it is an important part of our study of critical reasoning. In this chapter we will examine the infirmities of memory and learn some ways to guard against them.

### Contents

---

<b>7.1</b>	<b>Memory and Reasoning</b> . . . . .	<b>126</b>
<b>7.2</b>	<b>Stages in Memory</b> . . . . .	<b>126</b>
	7.2.1 Where Things can go Wrong . . . . .	128
<b>7.3</b>	<b>Encoding</b> . . . . .	<b>128</b>
<b>7.4</b>	<b>Storage</b> . . . . .	<b>129</b>
	7.4.1 Editing and Revising . . . . .	129
<b>7.5</b>	<b>Retrieval</b> . . . . .	<b>131</b>
	7.5.1 Context and Retrieval Cues . . . . .	134
	7.5.2 Schemas . . . . .	135
<b>7.6</b>	<b>Summary: Inference and Influences on Memory</b> . . . . .	<b>137</b>
<b>7.7</b>	<b>Chapter Exercises</b> . . . . .	<b>137</b>

---

## 7.1 Memory and Reasoning

The study of memory bears on our study of reasoning in three ways:

1. We often base the **premises** of our reasoning on what we think we remember. Since we usually trust our memories, such premises are usually thought to be especially secure.
2. Memory involves **inference**. This inference is influenced by our expectations, ways of labeling things, and even our biases, desires, and self-interest—the very things that so often lead to faulty reasoning.
3. Memory is susceptible to various sorts of errors, so we need critical reasoning to **evaluate** claims about what we remember.

Human memory is extremely impressive, but it can trip us up and some of its errors lead to errors in reasoning. How good are our memories? We have all used hundreds of phones. What letters go on which buttons (figure 7.1)?



Figure 7.1: What's on the Phone?

We have seen thousands of Lincoln-head pennies. Which image is correct (Figure 7.2 on the next page)?

If we can't remember the details of very familiar objects like phones and pennies, we may wonder how accurate people are about the details of things they see only briefly. How reliable, for example, is eyewitness testimony likely to be?

## 7.2 Stages in Memory

Remember your old view about memory? Well forget it

Most people think of memory as a storage device. It's like an information bank: we perceive something, store the information away in memory, and withdraw it later. We sometimes forget, to be sure, but when we don't forget, our memories are pretty




Figure 7.2: Which Penny is Right?

reliable. On this view memory is passive; it is a record of things we experienced or learned. But this view is wrong. Perception, as we saw earlier, doesn't work like a video camera. And memory doesn't work like a video tape. Memory is *active*, and it involves reconstruction on our part. What we remember is jointly determined by the information that does get stored in our brains *and* by our reconstruction of it when we remember something.

Memory is not a single, unitary process or system. In the 1960s, psychologists thought that memory consisted of short-term memory, which fades very rapidly, and long-term memory, which is more permanent.

Nowadays, they more frequently draw a distinction between *working memory*, which holds a small amount of information for a short period of time (like the phone number we keep repeating to ourselves as we scramble for the phone), and *long-term memory*. But there are increasing signs that working memory and long-term memory each consists of further subsystems. We don't need to worry much about this, however, and a simple three-part division will serve our purposes here.

We can think of memory as involving the following three stages:

1. **Encoding:** occurs when we perceive something
2. **Storage:** which often involves
  - (a) Elaboration (adding information to the memory)
  - (b) Revision ("rewriting" the memory)

### 3. Retrieval

- (a) Recall
- (b) Recognition

A crude flow chart of this is shown in Figure 7.3



Figure 7.3: Stages in Memory

#### 7.2.1 Where Things can go Wrong

The *output* of memory—the thing we actually think we remember—is often different, sometimes dramatically different, from the *input*, and errors can creep in at all three stages in the process. This isn't to say that our memories are wildly inaccurate. Creatures with extremely unreliable memories couldn't survive for long. Still, important errors often occur.

This is easy to overlook, because we often fail to notice when our memories are inaccurate. After all, our memories usually do seem very accurate *to us*. Moreover, the details often don't much matter, so we don't notice when they are wrong. Finally, it is often difficult, or even impossible, to check a memory against what really happened.

### 7.3 Encoding

We encode information when we perceive something, and if we misperceive something, then our memory of it will likely be distorted. Errors in the input usually lead to errors in the output. In the chapter on perception we saw a number of ways in which perception—including perception in the extended sense that encompasses our emotions and feelings—can be mistaken, so inaccuracies and biases can be encoded at the very beginning of the memory process.

For example, many of the spectators at the Princeton-Dartmouth game (p. 79) perceived things in a biased way, so it is little wonder that later they didn't have accurate memories of what occurred. But the relationship between perception and memory is a two-way street. Perception is the input for memory. But memory provides the basis for our perceptual set—what we expect to see is determined by our memories—so it in turn influences perception.

## 7.4 Storage

The chapter on perception shows that we sometimes misencode information, but once information is stored in the brain it might seem safe. Even here, however, our memories are active, and over time we unconsciously elaborate and revise the information we have stored. Since this occurs outside the realm of consciousness, it is sometimes difficult to determine whether the revisions occur during storage or during retrieval (for our purposes it won't usually matter which is involved), but in many of the examples in this section errors pretty clearly occur during the storage phase.

### 7.4.1 Editing and Revising

A central theme of this book is that we have a strong need to make sense of our world, to understand why the things that matter to us (including other people and ourselves) behave as they do. This drive for explanation and understanding is so strong that it sometimes leads us to see patterns and reasons even where they don't exist and to construct explanations even when we don't have enough evidence to warrant them. It can lead us to fill in gaps in our memories even when we have little objective basis for doing so. This will be clearer if we consider several examples that illustrate the varied ways we do this.

#### The Ants

Subjects in an experiment heard a story that contained sentences like:

- The ants ate the jelly.
- The ants were in the kitchen.

Later they were asked to identify the sentences they had heard. Most thought they remembered

- The ants ate the jelly in the kitchen.

But this sentence wasn't in the story. What happened?

The subjects had, automatically, filled in gaps based on what they knew made sense. They didn't store what they had literally heard, but an organized, meaningful version of the story. This filling in of gaps is a type of *inductive inference*. It is a way of *updating* the information stored in our heads. In this case the subjects had some stored information and then inferred things that seemed to follow from it. For example, they inferred that the ants ate the jelly in the kitchen. But this wasn't a conscious inference; they genuinely thought they remembered hearing the sentence.

### The Graduate Student's Office

In another study undergraduates were asked to wait in a graduate student's office. Later they were asked what was in the office, and most of them mentioned books. In fact there weren't any books in the office. What happened? The subjects' memories had added a detail, based on the subjects' expectations about graduate students' offices.

### The Dictator

In another study people heard a fictitious story about a dictator. In one version he was called 'Gerald Martin'; in another he was called 'Adolph Hitler'. The story didn't mention Jews. Many students in the group who heard the Hitler version of the story thought it contained the sentence

- He hated the Jews.

Students who heard the other version did not. What happened? Students in the first group filled in a detail based on what they knew about Hitler. They drew and inference, unconsciously, based on common knowledge.

### The Labels

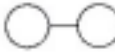











	Group 1		Group 2	
Figure shown to subjects	Label given	Figure drawn by subjects	Label given	Figure drawn by subjects
	<b>Eyeglasses</b>		<b>Dumbbell</b>	
	<b>Hourglass</b>		<b>Table</b>	
	<b>Seven</b>		<b>Four</b>	
	<b>Gun</b>		<b>Broom</b>	

Figure 7.4: Classification and Memory



In yet another study, people were shown several fuzzily drawn figures (figure 7.4 on the preceding page). Half of the people were shown the figures with one set of labels; the other half were shown the very same figures, but with different labels. For example, a figure that was labeled as a barbell for the first group was labeled as a pair of glasses for the second. The people were later asked to draw the figures they had seen. What do you think happened? As you might have predicted, the pictures they drew were heavily influenced by the labels they had seen. What they remembered seeing was partly determined by the way in which they had labeled or classified it.

### Betty K.

In our next example, subjects in a study heard a fairly neutral description of the early life of Betty K. The story contained sentences like

- Although she never had a steady boyfriend in high school, she did go on dates.

Later, half of the subjects were told that Betty later became a lesbian while the other half were told that she got married. The first group was more likely to remember “She never had a steady boyfriend” rather than “she did go on dates.” In the second group the results were reversed. What happened?

### The Lecture

A group of people attended a lecture. Some of them later read an *inaccurate* press report about it. Those who read this report tended to remember the lecture as it was described, even though the description was inaccurate. They unconsciously edited their memories in light of the report.

These five examples suggest a moral. Memory is not passive. It involves active reconstruction of things in an effort to make as much sense of them as we can. This reconstruction is influenced by our expectations and what we know. In the next section we will see that it can be affected by other things as well.

## 7.5 Retrieval

There are two forms of retrieval. In **recall** we actively remember a fact, name, etc. The example of the phone buttons requires you to recall which letters go with which numbers. By contrast, in **recognition** we only have to recognize something when we perceive it. The Lincoln-head penny example doesn't require you to describe

or recall the face of a penny; it simply asks you to recognize the correct picture when you see it.

Although *retrieval* is a natural word for the elicitation of information from memory, *reconstruction* would often be more accurate. Retrieval is the *joint effect* of what is actually stored in the brain *and* of our present inferences about it. You can begin to see this if you try to remember the things you did yesterday and the order in which you did them. Yesterday's events do not pop up in memory, one by one, in the right order. You have to do some reasoning to see what makes sense. It might go something like this:

Well, let's see, at noon I drove to Wendy's, but since I stopped by Homeland on the way I must have gone there before Wendy's. Then I went to the bank. Hmm . . . . No that can't be right. That doesn't make sense, since I was broke and I had to go get money from the bank to pay for my moon pie and fries. So I guess I went to the bank between going to Homeland and Wendy's. . . .

The way we reconstruct things in memory is influenced—sometimes dramatically—by the *context* in which we remember. One way context affects memory is by providing retrieval cues. **Retrieval cues** are features of the situation that help us retrieve information from memory. For example, if you are trying to recall someone's name, picturing them or recalling other information about them often helps you to remember.

Memory of an event occurs (by definition) *after* that event, and many things going on at the later time affect what we remember, how we remember it, and the way that we organize it into a meaningful pattern. Not only do we fill in gaps to help make sense of the earlier event; our memory of an earlier event is also colored by our attempts to make sense of the present.

Many features of a context can influence our reconstruction of the past. These include our current beliefs and attitudes, emotions and moods, expectations and set, motivations and goals (including the goals to look good and maintain self-esteem), the way questions are worded, and other peoples' suggestions. We will now examine the ways such factors can influence our memories.

### **Current Attitudes and Beliefs**

We tend to remember our earlier beliefs, opinions, attitudes and even our behavior as being more like our current beliefs and attitudes than they actually were. Greg Markus conducted a ten-year study of changes in people's political attitudes over time. In 1973 he surveyed a group of graduating high-school students along with many of their parents. He asked them about their attitudes toward the legalization

of marijuana, women's rights, affirmative action programs, equality for women, and several other social issues. Ten years later he asked the same people (i) what their *current* attitudes on these issues were and (ii) what their *earlier* attitudes, in 1973, had been. Both the students' and the parents' memories of their earlier attitudes were *much* closer to their current attitudes than to the attitudes they had actually expressed back in 1972.

In another study people were asked to report on their political views in 1972. Four years later they were asked what their current views were and what their earlier views had been. Many people's views hadn't changed, and 96% of the people in this group (correctly) reported that their views had remained constant. But some people's views had changed, and 91% of them (incorrectly) reported that their views had *not* changed.

People sometimes also remember their earlier behavior as being more in line with their current views and behavior than it actually is. Linda Collins and her coworkers asked high-school students about their use of tobacco and alcohol. Two and a half years later they asked them (i) what their *current* patterns of use were and (ii) what their *earlier* pattern of use, two and a half years earlier, had been. Their memories of their earlier pattern of use were closer to their current pattern than to the pattern they had reported earlier.

These results may explain why each generation of parents and teachers wonder why the current generation seems to be going to hell in a handbasket: "Why can't today's teenagers be more like *we* were when we were young?" Parents and teachers may be comparing their remembered version of their past (which is much more like their current views than their own past actually was) with today's generation, rather than comparing how things really were in the past with today's generation.

The effects discussed here are relatively modest, and people often do accurately recall their earlier views. But there is a definite tendency to see our earlier beliefs and attitudes as more like our current beliefs and attitudes than they actually were. This fosters the view that our beliefs and attitudes are more stable and consistent over time than they actually are. This can lead us to suppose that our *future* beliefs and attitudes will be more like our present ones than they will turn out to be. To the extent that this happens, we have an inaccurate picture of ourselves.

### **Current Moods and Emotions**

Cognition and emotion—thought and feeling—are more intertwined than we sometimes suppose, and our moods and emotions can affect memory. Although the evidence is cloudy, there is some evidence that people who learn material in one mood recall it more easily when they are in that mood. And studies of actual patients over a several year period showed that when people are sad or depressed they tend to

remember more negative things. For example, they are more likely to remember their parents as unsupportive, rejecting, even unloving, than people who aren't depressed. This raises the question whether people are depressed because they had a bad childhood *or* whether they tend to remember having a bad childhood because they are depressed (it could be a bit of both).

### 7.5.1 Context and Retrieval Cues

It is often easier to remember something if we are in the context where we experienced it. This is called **context-dependent retrieval**. Being back in the original context jogs the memory by providing more *retrieval cues*. For example, you would probably find it easier to remember names of last year's acquaintances if you walked back through your old dormitory; it's full of cues that would help you remember the people who lived there. Or suppose that you are in your kitchen and think of something that you need to do on the way to campus. You walk into the hall and can't remember what it was. Often going back to the kitchen helps you recall; it contains cues that help you remember what you forgot.

The importance of context shows up over and over. For example, students do better when they are tested in the room in which they learned the material. And smells are particularly powerful at evoking memories that are associated with them; they provide a cue that can awaken memories that are hard to access in other ways.

It is also easier to remember something if we are in the same physiological state that we were in when we learned it. Here the context is physiological, inside our skins, and our own internal states provide a retrieval cue. This is called **state-dependent retrieval**. For example, if you learned something after several drinks or cups of coffee, it will probably be easier to remember under those conditions.

### Framing Effects: The Collision

When someone asks us to remember something, the way they word or frame their request can influence what we remember. Half the people in a group were asked "How frequently do you have headaches?" and the other half were asked "If you occasionally have headaches, how often?" The average response of the first group was 2.2 headaches a week while that of the second group was 0.7 headaches a week. Similarly, it has been found that if you survey the people coming out of a movie and ask half of them "How long was the movie?" and the other half "How short was the movie?" those asked the first question will think the movie was longer.

In a study having more obvious real-life implications Elizabeth Loftus and her coworkers asked subjects to watch a film of a traffic accident. Later they were

asked:

How fast were the cars going when they \_\_\_\_\_ each other?

The blanks were filled in with different verbs for different groups of subjects. When people were asked how fast they had been going when they smashed each other subjects remembered them going faster than when they were asked how fast the cars were going when they contacted each other. The results were:

1. smashed: 40.8/mph
2. hit: 34.0/mph
3. contacted: 30.8/mph

They were also more likely to remember seeing broken glass at the scene, even though none was present, when the collision was described in the more violent terms. Here the way the experimenter worded things affected what people remembered. If such small changes of wording can produce such dramatic effects, we must wonder what effects leading questions from a skillful lawyer, hypnotist, or therapist might have.

### 7.5.2 Schemas

We can tie some of these examples together with the notion of a schema. Our beliefs about the world in general also play a role in our construction of memories. Consider the sentence:

- Wilbur was annoyed when he discovered he had left the mustard out of the basket.

What is the setting? Why should the mustard have been in the basket? Where is Wilbur likely to be when he discovers the mustard isn't there? Someone from another culture might have trouble answering these questions, but you saw straightaway that Wilbur has gone on a picnic and that he left the mustard out of the picnic basket. That was easy—but *how* did you know this?

There is now considerable evidence that we have well-organized packets of generic knowledge about many things, including picnics, graduate student offices, classrooms, visits to restaurants, first dates, and so on. These packages of information are called **schemas**. We won't worry about the exact nature of schemas, which isn't well understood in any case, but the basic idea will be useful.

Most of us have a packet of information about the typical picnic, a picnic schema. In the typical picnic people we pack food in a picnic basket, take along ketchup and mustard, eat outside, and so on. We can have picnics without any of these features, but such things are part of our picture of a typical picnic.

Schemas are very useful because they help us organize our knowledge and automatically fill in many details. A little information may activate the schema, and then we use the generic knowledge in it to quickly draw further *inferences* about the situation or thing. For example, mention of a basket and mustard activate our picnic schema, and we can then use it to draw inferences about what Wilbur is up to. Similarly, your schema for a graduate student office probably includes having books in it, so it is natural to infer that it does.

Schemas enable us to form accurate expectations about a situation on the basis of just a little information about it. These expectations may be wrong, but we will often be surprised if they are. For example, our schema of a classroom includes having a roof, and if you walked into a classroom and found no roof, you would be surprised.

Schemas figure in memory in the following way. If you remember a few fragments of experience that activate a schema you then have a tendency to remember other things that are included in that schema. The knowledge in the schema helps you to fill in the gaps. Often this filling in is accurate. Most graduate student offices *do* contain books. Again, it is part of many people's schemas of classrooms that they have fluorescent lighting. It turns out that many people think they remember that a given classroom had fluorescent lighting, even if they didn't notice the lighting. In most cases classrooms *do* have such lighting, so more often than not this gap in memory would be filled in accurately. But we will be mistaken when we are asked about a classroom with some other sort of lighting.

Many schemas are accurate. Graduate student offices usually contain books, classrooms usually have fluorescent lighting, people on a first date are usually on their best behavior. When we base our inferences, including the inferences involved in memory, on such schemas we will often be right. But sometimes we will be wrong: some student offices don't contain books; some people on first dates are very nervous and you see them at the worst.

### **Stereotypes**

Not all schemas are accurate. **Stereotypes** are schemas, mental pictures we have of clusters of traits and characteristics that we think go together. Most of us have various racial, ethnic, and gender stereotypes. Many of these are inaccurate, and they can lead us to perceive and remember and infer things in a distorted way. For example, you may have a stereotype about the typical New Yorker that includes being rude and pushy. If so, you are more likely to predict that a given New Yorker will be pushy, more likely to interpret a New Yorker's behavior as pushy, and more likely to remember the behavior as pushy. People who heard the story about Betty K (page 131) probably edited their memories in light of their stereotype of lesbians.

## 7.6 Summary: Inference and Influences on Memory

The examples we encountered show that memory involves a good deal of reconstruction or inference, and this reconstruction is highly sensitive to context. What we remember can be influenced by:

1. Obvious inferences (the ants)
2. Common knowledge (Hitler story)
3. Expectations (graduate student office)
4. Labels and concepts (labeled figures)
5. Schema (Betty K and stereotype of lesbians)
6. Subsequent information (the lecture)
7. Current attitudes and beliefs (attitudes towards drugs)
8. Current moods and emotions (attitudes of depressed people)
9. Nature of retrieval cues, e.g., subsequent framing (the collision)

In short, there is a good deal of evidence that when we remember something we are engaged in a sort of inference that moves from information stored in our brain and the features of the situation in which we remember to a conclusion about what we originally saw, heard, learned. This isn't a defect of memory. Indeed, it shows some intelligence to automatically try to make sense of things and fill in gaps and focus on essentials rather than on irrelevant details (like the exact wording of the sentences about the ants). It's just that sometimes the inferences lead us astray.

Many of the things on this list can also influence reasoning, and memory is susceptible to many of the same kinds of errors that reasoning and inference are. Because of this, memory can be critically evaluated just like any other source of information can. In the next chapter we will examine common errors in memory and learn about some ways to avoid them.<sup>1</sup>

## 7.7 Chapter Exercises

1. The way we word questions can affect the way people remember things. Give two examples of different ways of wording the same question that might elicit different memories. How could you test whether your questions really did this?
2. Jack was walking home late last night after a few too many drinks. He sees someone who may just have broken into his neighbor's house, but can't remember much about the burglar. Under what conditions might he be more likely to remember?

---

<sup>1</sup>References on memory will be found at the end of the next chapter.

3. Most of us have trouble with the penny identification exercise. This is nothing to be worried about; people who spent much time memorizing the details of such things need to get a life. Typically we only need to know enough to recognize pennies when we see them in the real world. But when the situation changes, things that were unimportant may become important. For example, no one cared much about the appearance of quarters until the Susan B. Anthony \$1 coin was introduced. It quickly fell out of favor because it was easily confused with a quarter. Given another example where remembering the details about something didn't matter until the situation changed. What is the moral of such examples?

#### **Answers to Selected Exercises**

1. There is typically little reason to remember much about how someone we see at a glance really looks. But if we learn they are the kidnapper, it becomes relevant. But think up an example of your own. One important moral of such cases is that what is important and relevant often depends on context.



# Chapter 8

## Memory II: Pitfalls and Remedies

**Overview:** In the previous chapter we learned about several basic features of memory and a few of the ways they can trip us up. In this chapter we will learn about additional pitfalls and some remedies for them.

### Contents

---

<b>8.1</b>	<b>Misattribution of Source</b>	<b>140</b>
<b>8.2</b>	<b>The Power of Suggestion and the Misinformation Effect</b>	<b>140</b>
<b>8.3</b>	<b>Confidence and Accuracy</b>	<b>141</b>
8.3.1	Flashbulb Memories	141
<b>8.4</b>	<b>False Memories</b>	<b>142</b>
8.4.1	Motivated Misremembering	143
8.4.2	Childhood Trauma and False-Memory Syndrome	143
<b>8.5</b>	<b>Belief Perseveration</b>	<b>144</b>
<b>8.6</b>	<b>Hindsight Bias</b>	<b>145</b>
<b>8.7</b>	<b>Inert Knowledge</b>	<b>145</b>
<b>8.8</b>	<b>Eyewitness Testimony</b>	<b>146</b>
<b>8.9</b>	<b>Primacy and Recency Effects</b>	<b>147</b>
8.9.1	The Primacy Effect	147
8.9.2	The Recency Effect	147
<b>8.10</b>	<b>Collective Memory</b>	<b>148</b>
<b>8.11</b>	<b>Remedies</b>	<b>149</b>
8.11.1	Safeguards	149
8.11.2	Ways to Improve Memory	149

8.12 Chapter Summary . . . . .	151
8.13 Chapter Exercises . . . . .	152
8.14 Appendix: Different Memory Systems and False Memories	154

## 8.1 Misattribution of Source (“Source Amnesia”)

**Misattribution of source:**  
forgetting where we first  
learned something

Often we remember something accurately, but we forget what the source was. We may even form quite mistaken beliefs about the source. We are wrong about where we learned it, when, and who we learned it from. For example, Ronald Reagan was fond of telling a story about a World War II gunner whose plane was severely hit by enemy fire. His seat ejection device malfunctioned, and the commander said “Never mind, son, we’ll ride down together.” The commander, Reagan said, was awarded the Congressional Medal of Honor for this heroic act.

It turned out that no medal had ever been awarded for this action, but that the scene had occurred in the 1944 movie *A Wing and a Prayer*. Reagan correctly remembered the story, but he was wrong about its source.

In extreme cases we may not know whether the source of a current “memory” is an earlier event or something we merely imagined. And in some cases, where we can’t actually remember a source, we (unconsciously) invent one. The hypnotist tells me that when I come out of my trance I will crawl around on the floor when he snaps his fingers. I later react to his cue when he snaps his fingers and someone in the audience asks me why I am crawling around on the floor. I will very likely make up an explanation right there on the spot—I’m looking for a pen that I dropped—and what’s more *I’ll* believe it. In this case I am the one taken in by the story I invent.

It is easy to laugh at some misattributions of sources, but we all make mistakes of this sort. Often our errors are harmless, but sometimes they aren’t. It is likely that some cases of plagiarism involve source amnesia. An author gets an idea from reading someone else, then later forgets that he read it and thinks that the idea is his own. Or we may think we learned something from a reliable source when in fact we got it from someone unreliable, we will be more confident in our belief than we should be.

## 8.2 The Misinformation Effect

The psychologist Elizabeth Loftus had subjects view a simulated automobile accident at an intersection near a stop sign. Later, the experimenters suggested to one

group that the sign was a yield sign. Still later, when subjects in this group were asked about the sign, many thought that they *had* seen a yield sign (subjects who had not heard this false suggestion were much more accurate about the sign).

In another study Loftus showed subjects a videotape in which eight demonstrators burst into a classroom and disrupted things. Afterwards she asked half of the subjects whether the leader of the twelve demonstrators was male and she asked the other half whether the leader of the four demonstrators was male. A week later she asked the two groups how many demonstrators there were. The average response of those who had been asked about four demonstrators was 6.4; the average for the half who had been asked about twelve was almost 9.

Subjects in these studies were victims of the **misinformation effect**. When we are exposed to subtle, even barely noticeable misinformation, the misinformation often influences our memories. This occurred when subjects revised their memory of a lecture in light of a later report that they read about it. We could think of leading questions as a form of subtle suggestion, so the subjects who were shown the film of a collision were also victims of this effect. Clearly our ability to reason critically and accurately will be impaired when our beliefs about the past are distorted by other people's manipulations, whether intentional or not, of our memories.

**Misinformation effect:**  
after exposure to subtle misinformation, people often misremember

## 8.3 Confidence and Accuracy

### 8.3.1 Flashbulb Memories

Memories of some highly emotional moments seem particularly vivid and indelible; we are very certain that we accurately remember the details surrounding them. Most of you will have a clear and confident answer to the first question (and if you are old enough, to the other two as well).

1. Where were you when you learned about the bombing of the Alfred E. Murrah Federal Building in Oklahoma City (April, 1995)? How did you hear about it? What else was going on around you then?
2. Where you were when you learned that the Challenger spacecraft had exploded (January 1986)? How did you hear about it? What else was going on around you then?
3. Where you were when you learned about Kennedy's assassination (November 22, 1963)? How did you hear about it? What else was going on around you then?

The events are so dramatic that it feels like a mental flashbulb went off, freezing a snapshot of things indelibly in our minds.

How likely is it that your memory about how you learned of the bombing are mistaken? How would you feel if they turned out to be dramatically wrong? The day after the Challenger disaster the psychologists Ulric Neisser and N. Harsch asked a large group of undergraduates to write down where they were when they learned about it, who they heard it from, and so on.

Two and a half years later, these people were again interviewed about the setting in which they learned about the explosion. The accounts of over one third of the people were quite wrong, and another third were partly wrong. When they were shown the statements they wrote right after the explosion, many of the people were very distressed. Most of them preferred their recent account to the original one; they thought it more likely that they had been mistaken the day after the disaster than two and a half years later.

We are so confident of these memories that the possibility they could be distorted disturbs us. But it turns out that there is not a high correlation between the vividness of a person's memory and its accuracy. Nor is there a high correlation between a person's confidence in a memory its accuracy. Confidence and vividness are very fallible indicators of the accuracy of a memory.

## 8.4 False Memories

**False memory:** thinking we remember something that in fact never happened

It is one thing to misremember the details of something; it is quite another to think you remember something that never happened at all. Fortunately the former is far more common than the latter, but the latter does occur. For example, there are a number of people who think they remember being abducted by aliens or having a past life or (closer to home) events from their childhood that never occurred. A **false memory** is one that is very inaccurate; it may even be a "memory" of something that didn't happen at all.

Corroboration by another person is one of the most potent causes of false memories. Many other things, even imagining something or recounting stories about it, can lead to false memories. For example, Maryanne Garry asked students about various kinds of events that occurred when they were children.

Two weeks later she had students vividly imagine that they had experienced various events occurring as children, e.g., that they hit a window, broke the glass, and cut their hand. Some of the students actually came to *believe* that they really had experienced such events, after *imagining* that they had done these things years before. One of the chief problems in assessing the accuracy of memories is that *false memories often feel just like accurate memories.*

### 8.4.1 Motivated Misremembering

Our desires and motives also sometimes lead us to misremember things. Most of us like to see ourselves in a good light, and so we are likely to remember things in a way that will protect our self-image and self-esteem. To see how common this is just remember some situation where two people who live together argue about something, e.g., who did their fair share of the housework or who is to blame for various problems. The two people probably had rather different memories about who did what, and each remembered things in a way that put them in a better light. Often both people's claims are sincere—they really think that their memories are accurate—but they can't both be right.

Freud thought that we used defense mechanisms to protect our view of ourselves. A **defense mechanism** is something we do (typically unconsciously) to keep from recognizing our actions, motives, or traits that might lower self-esteem or heighten anxiety. Some defense mechanisms involve distortions of memory. The most extreme case is repression, forgetting things that are unpleasant to remember or face. This is the opposite of a false memory. Rather than remembering something that didn't happen, we erase the memory of something that did. It isn't clear how often repression actually occurs, but it is clear that our memories are often self-serving. We reconstruct the past in a way that puts us in a good light.

### 8.4.2 Childhood Trauma and False-Memory Syndrome

In recent years a number of psychologists have argued that some childhood events (e.g., sexual molestation) may be so traumatic that people repress them. The memories would be so painful that the victims simply forget that they ever happened. But although they can't bring the experience into consciousness, its traces still linger in some form that leads to long-term problems like low self-esteem, depression, and sexual dysfunction.

Many people, often with the aid of therapy, have uncovered what they think are memories of childhood traumas like sexual abuse. How accurate are these memories? It is clear that children are sexually molested more often than society used to suppose, and in some cases very traumatic events are forgotten. So some of these memories are probably accurate.

Recently, however, a number of psychologists have argued that many of the reawakened memories are really false memories, and they urge that the victim is actually suffering from false-memory syndrome. **False-memory syndrome** is a pattern of feelings, emotions and thoughts based on distorted or completely false memories.

Those who believe that many of the alleged victims are really suffering from

**False-memory syndrome:** a pattern of feelings and thoughts based on false memories

this syndrome point out that although traumatic events are sometimes forgotten, they are usually remembered all too well. Furthermore, childhood memories of events that occurred before age three are very unreliable (the parts of the brain needed to store memories simply haven't developed enough before then). Later childhood memories are often accurate, but source misattribution and misinformation effects plague us all. Children have been shown to be particularly suggestible, and they often receive subtle—and sometimes not-so-subtle—suggestions from a parent or a therapist.

Suggestibility is a less serious threat in older patients, but they may be motivated to misremember. They may find the idea that they were abused as children a convenient way to place the blame on someone else for their current problems on some one else.

There is a great deal of debate about the extent to which people's reports of reawakened memories of childhood traumas are false memories. The issue raises a terrible dilemma that has torn many families apart. Childhood abuse does occur, and when it has it is important for the victim to try to deal with it in an open and honest way. On the other hand, the effects of false charges of molestation by family members are catastrophic. The problem is that we usually can't check a person's childhood to see whether the reported memory is accurate or not.

## 8.5 Belief Perseveration

There are additional problems with memory that pose obstacles to clear and careful thinking, and we will examine several of them in the remainder of this chapter.

**Belief perseveration:**  
tendency to continue to believe something even after we get evidence it isn't so

Many studies, and much ordinary experience, show that we have a tendency to retain a belief even after our original reasons for thinking it true have been undermined. Such beliefs are so thoroughly entrenched that they are impervious to evidence that would discredit them. This phenomenon is known by the ugly name of **belief perseveration**. Once something gets into memory, whether it is accurate or not, it can be difficult to get out.

Belief perseveration occurs in many psychology experiments. If the subjects in certain psychology experiments knew the true purposes of the studies they might act in ways that would undermine the experiments, so they are often given a false cover stories. For example, people in an experiment about conformity might be told that they are in a study about perception. That way subjects are more apt to behave as they normally would. Once the experiment is over, the experimenter is required to "debrief" the subject, to explain what the experiment was really about. But it has been found repeatedly that even when the true purpose of the experiment is explained in detail, many subjects persist in thinking that the earlier account, the

false cover story, was right. They are victims of belief perseverance.

Many of our stereotypes are also resistant to change, even after we meet numerous people in a group and discover that most of them don't fit our stereotype. We will examine some of the reasons our beliefs are resistant to change in later chapters.

## 8.6 Hindsight Bias

When we learn some fact or the outcome of some event, we have a strong tendency to think that we would have predicted it before hand. This "I knew it all along" syndrome is called **hindsight bias**. Hindsight bias has been found in elections, medical diagnoses, sporting events, and many other settings.. Hindsight bias tends to confirm our view that we are right ("I knew what would happen") more often than we actually are. This makes it more difficult to correct our mistakes by learning from past errors, since we don't even notice them.

**Hindsight bias:** I knew it all along

### Remedies

Warning people of the dangers of hindsight bias has little effect. But we can reduce it by considering how past events might have turned out differently. Ask yourself what alternatives might have occurred and what things would have made it likely that they would have happened?

## 8.7 Inert Knowledge

We have a great deal of knowledge stored in memory that we can't access when we really need it. We know it, but we just don't think about it in cases where it actually applies. The philosopher Alfred North Whitehead called this **inert knowledge**.

**Inert knowledge:** knowledge we can't access when we need it

Most of you remember the Pythagorean theorem. It tells us how to calculate the length of any side of a right triangle if we know the length of the other two sides (if  $h$  is the length of the hypotenuse and  $a$  and  $b$  are the lengths of the other two sides, then  $h^2 = a^2 + b^2$ ). On two different occasions I have seen people engaged in minor construction (building a dog house and building a bookcase) who went to all sorts of trouble to figure how long various boards should be cut. They could easily have answered their questions by using the Pythagorean theorem, but it just didn't occur to them to do so. When I mentioned the theorem they remembered it, but they hadn't realized that it applied in the case at hand. They didn't "code" the situation as one where it was relevant, so their knowledge of the formula remained inert, dormant, unactivated.

The chief problem with many courses, including critical reasoning courses, is that the knowledge you acquire in them will remain inert. When you take a class you usually remember things when you need them on a test. But the material is useless unless you can also apply it outside of class. It is surprising, and depressing, how difficult this is. Some of you will become teachers, and this will be a problem you will constantly face. It helps if you think of examples of things, e.g., belief perseveration, from your own life and if you watch out for them outside of class. It also helps if you learn to recognize *cues* that signal the relevance of something you have learned (like the rules for calculating probabilities, which we learn below) to a given problem.

## 8.8 Eyewitness Testimony

Many crimes would never be solved without eyewitness testimony. But how good is it? Study after study shows that people in general, and jurors in particular, put great confidence in the testimony of eyewitnesses. It has also been found that the more confident a witness sounds, the more persuasive they are.

But it has also been found witnesses who seem the most confident are by no means always the most accurate. Indeed, eyewitnesses, whether confident or not, are often mistaken. Even when they make every attempt to be honest and conscientious, as most do, their memories are subject to all of the infirmities discussed earlier in this chapter. Indeed, many studies show that the descriptions of eyewitnesses are often dramatically wrong, and many innocent people have been convicted on the basis of well-meaning, but inaccurate, eye-witness testimony.

By now the fallibility of eyewitnesses shouldn't seem that surprising. After all, they often get only a quick glimpse of the perpetrator. In the case of violent crimes, they are likely to feel stress and fear, which degrade the ability to remember details. There are further problems with police line-ups. Most witnesses believe that the guilty person is in the line-up, so they often select the person who looks most like the person they saw. Furthermore, police prompting or leading questions can lead to misinformation effects.

You have now learned enough about memory that you should be able to think of various things that would enhance the reliability of eyewitnesses and, indeed, of anyone else who is trying to remember some event. What might you do? Could you employ retrieval cues? What things should the questioner avoid doing? We will return to these questions below, but you should try to answer them now.



## 8.9 Primacy and Recency Effects

### 8.9.1 The Primacy Effect

John is envious, stubborn, critical, impulsive, industrious, and intelligent. In general, how emotional do you think John is? (Pick one).

**Primacy effect:** early items influence us more than later ones

Not emotional 1 2 3 4 5 6 7 8 9 Extremely emotional

Some memory effects depend on time, on the temporal context. In particular, early items in such a list like this influence our impressions and inferences about someone much more strongly than later items do. This stronger influence of earlier items or situations is called the primacy effect. The **primacy effect** means that events or features appearing early in a series are easier to remember than later ones.

Other things being equal, first impressions (and to a lesser extent second and third impressions) have a stronger impact on people than later impressions. There are many situations where this matters, including the first impression you make on a date or at a job interview. The first sentence, paragraph and page of a paper make a first impression on its readers, and if it is bad, they aren't likely to like the paper as a whole. The first quiz or exam you take in a course might also influence the professor's evaluation of your work generally.

One plausible explanation for the primacy effect is that we have a limited amount of time and attention, so we can't constantly monitor everything and update our views about it constantly. If this is correct, the effect may involve attention and perception as much as memory, but it shows up when we think back over the series of things. By giving more weight to early impressions than to later ones, we rely on a *biased sample*. This leads us to draw inferences based on inductively weak reasons. Such inferences are flawed. But if we are aware of this tendency we can try to avoid it in ourselves. We can also try to create the best first impression that we can.

### 8.9.2 The Recency Effect

Although the first things in a series tend to be given more weight than later one, things at or near the end of the series may also be given more weight than average. This seems intuitive: things near the end are fresher in our minds.

**Recency effect:** later items also influence us more than ones in the middle

Which effect is stronger? There is some evidence that when presentations are involved, there can be a strong recency effect. But when we judge other people, the primacy effect is much stronger. But even this depends on the context. Early in their relationship, Sarah's first impression of Wilbur loomed larger than later

impressions. But once they have been married for thirty years, the first impression won't matter as much as later impressions.

In a detailed study, Miller and Campbell edited court transcripts in a case seeking damages due to a defective vaporizer. All the material supporting the plaintiff was placed together in one long block, and all the material supporting the other side was placed in another long block. The existence of primacy and recency effects depended on delays in the process.

1. If people heard each side's case and then made a judgment about them, neither a primacy nor a recency effect was observed.
2. If people heard both messages back-to-back, then waited a week to make a judgment, there was a primacy effect (the view presented first did better).
3. If people heard one message, waited a week to hear the second, then immediately made a judgment, there was a recency effect (the view presented last did better).

What might explain this? When there was a week between presentations, people remembered much more about the side presented last. This seems to explain the recency effect in these conditions. What might explain the primacy effect where it occurs? The study suggests a bit of advice. Speak first if the other side will speak right afterwards and there is a delay in the response to the presentations. But speak last if there will be some time between the two presentations and the response to the presentations will come right after the second.

## 8.10 Collective Memory

**Collective memory:** how an entire group remembers something (often inaccurately)

Individuals have memories that are stored in their brains. Societies and cultures, have a sort of collective memory that is embodied in their beliefs and legends and stories about the past. Social scientists have found that collective memories change over time and such memories are often quite different from the original events that gave rise to them.

Sometimes people in power, particularly in totalitarian societies, set out to revise collective memory. There are many techniques for doing this, including rewriting textbooks, constantly repeating the rewritten version of history, and forbidding discussion of what actually happened. Some of the things that lead to distorted memories in individuals are different from those that lead to distortions in collective memory. But can you think of any similarities?

## 8.11 Remedies

### Hypnosis is not the Answer

We often hear that people can recall things under hypnosis that they couldn't otherwise remember, and to some extent this is true. But hypnosis does not provide a magic route into memory. In fact, people under hypnosis are particularly susceptible to misinformation effects, often stemming from leading questions from the hypnotist. They are also susceptible to misattribution of source, since an accomplished hypnotist can often get them to believe that they really remember something which the hypnotist in fact suggested to them while they were hypnotized.

Indeed, hypnosis is often an effective technique for implanting false memories. Hypnotists can induce their subjects to "recall" all sorts of outlandish things while under hypnosis, including alien abductions and events from "previous lives."

### 8.11.1 Safeguards

We can put the things we have learned together by asking which things make memories more accurate and which things make them less so. The lessons are completely general, but to make this more concrete, think about eyewitnesses.

What would help make eyewitnesses *more* accurate? From what we have learned, it would be useful to encourage them to use *retrieval cues* by asking the witness to visualize the crime scene, recall the weather and time of day, remember their mood, sounds, the obvious things that they saw, and so on.

What would make eyewitnesses *less* accurate? We should avoid asking leading questions or making even vague suggestions that could lead to source misattribution or misinformation effects. It would be better to encourage the witness to recount things without any interruption and noting every detail, however trivial, they can. We should only ask questions after they have completed their story.

These techniques do improve the reliability of witnesses, in some studies by as much as 50%. The idea is simply to use the things that improve memory, (e.g., retrieval cues), while steering clear of the more easily avoided phenomena, (e.g., leading questions), that lead to mistakes. The lessons here are perfectly general. They apply to a therapist trying to find out about a client's childhood, in daily life when we wonder if someone else's memories are accurate, and to ourselves when we are trying to remember some fact or event.

### 8.11.2 Ways to Improve Memory

There are many techniques for improving your memory. We will focus on those that are especially relevant to college students, namely doing a better job of re-

membering what you learn in your courses, but many of the points apply to a much wider range of settings.

### **Be Active**

It is almost impossible to remember (or even to fully understand) material if you are passive. If you just sit back and listen, you will retain very little of what you hear.

### **Integrate the material with things you already know**

We remember what we understand

An important way to be actively involved with material is to organize it in a way that allows you to fit it into a pattern that connects with things you already know. We are much better at remembering things we understand. So you need to integrate new material with your current knowledge. If you have integrated the material with the rest of the things that you know, you will see a pattern, connections to lots of other things, and that will help you remember and apply it.

In order to integrate new information with things you already know, you need to put things in your own words, think of examples from your own life, ask how the principles would apply outside of class. Otherwise the material will simply seem to be a series of isolated, unrelated facts that don't add up to anything.

Here's an example from earlier in this class. Most of you can repeat the definition of a *valid argument*. But unless you understand why validity is so important and learn to distinguish valid from invalid arguments, you won't really understand the definition or be able to integrate it with other things you have learned. In fact, you won't even be able to remember the simple definition for very long.

### **Lectures**

You will play a more active role in your own learning if you sit near the front of the classroom and get involved. You should take notes; this not only gives you a record to study later, but it reduces passivity. But the notes should be brief, in an outline style. Write them as if you were writing newspaper headlines: pack as much information as you can into the fewest words. You should also put things down in your own words. This requires more mental effort, but this effort will make it easier to understand and remember the material.

Most forgetting takes place soon after learning. So it is very useful to quickly review your notes soon after a lecture. It is also better to spend a few minutes reviewing it several times than to spend a lot of time reviewing it once.

### Tapes aren't much Good

Tape recording lectures is a very bad way to learn. It encourages passivity. You just sit back, tune out, and let the tape recorder do the work. But the tape can't distill and organize material, much less put things into your own words. You also lose out on the visual information, which is better for conveying certain concepts and principles than words are. Finally, the tapes of all the lectures between two exams will be 20 or 30 hours long, and the chances that you—or anyone—would actually listen to all of them are very slim. And even if you do, you will be so overwhelmed by 20 hours of tapes that you won't be able to remember much of what you hear.

### Reading

Quickly skim through section titles to get a sense of organization. Then read with your mind in gear; if something is unclear, try to understand it before going on to the next thing. Finally, after each section, pause and ask yourself what its main points were and try to think of ways the material is relevant in your life outside the classroom.

## 8.12 Chapter Summary

1. Memory involves a good deal of reconstruction and this reconstruction is highly sensitive to context.
2. What we remember is affected by our expectations, emotions, labels, and the context in which we remember it.
3. Hence, memory is susceptible to many of the same kinds of errors that reasoning and inference are.
4. Because of this, memory can be critically evaluated just like any other source of information can be.<sup>1</sup>

---

<sup>1</sup>Good, up-to-date, and reasonably accessible discussions of numerous aspects of memory distortion can be found in Daniel L. Schacter, ed., *Memory Distortion: How Minds, Brains, and Societies Reconstruct the Past*, Cambridge, Ma: Harvard University Press, 1995. The example of Betty K. is from M. Snyder and S. W. Uranowitz, "Reconstructing the Past: Some Cognitive Consequences of Person Perception," *Journal of Personality and Social Psychology*, 1978 (247); 941-950. A clear discussion of the influence of current views and moods on memory, together with references to some of the relevant literature, can be found in Robyn M. Dawes, *Rational Choice in an Uncertain World*, Harcourt Brace, 1988, Ch. 6. Elizabeth Loftus discusses her work in "The Incredible Eyewitness," *Psychology Today*, December 1975 and here work on false memories in "Creating False Memories," *Scientific American*, September 1997, 70-75; the biography here lists other useful sources. John D. Bransford, Jeffery Franks, Nancy Vye, and Robert Sherwood discuss the problem of inert knowledge in "New Approaches to Instruction: because Wisdom cant be Told," in S. Vosniadou, et al., eds., *Similarity and Analogical Reasoning*, 1989; 470-497. For a recent discussion of split brains and their

### 8.13 Chapter Exercises

1. If a professor is trying to learn the names of her students, which names do you think she'll find it easiest to learn? Would names at the beginning be easiest? The middle? The end? Why? How would you test a hypothesis about this?
2. If you are interviewing for a job, would you like your interview to come near the beginning? The middle? The end? Why? How would you test a hypothesis about this?
3. I have a very vivid memory that I heard about Kennedy's assassination during lunch break while I was standing out in front of my school. At least this is how it seems. But after learning about the tricks memory can play on us, I'm not so sure this is accurate. What might I do to discover whether this is accurate or not?
4. Suppose that you were teaching a course in critical reasoning. What would you do to combat the problem of inert knowledge?
  1. Give a specific, detailed (half a typed page) example of an assignment on the concept of deductive validity that is designed to give students a working (rather than merely inert) grasp of the concept.
  2. Give a specific, detailed (half a typed page) example of an assignment about some topic in this chapter that is designed to give students a working (rather than merely inert) grasp of the concept.
5. Think of a case where you held some belief long after the evidence suggested that you should abandon it. Write a paragraph describing this situation; include a discussion of things that might have led to the perseverance of this belief.
6. One problem in assessing reported memories of childhood traumas is that we usually can't check a person's childhood to see whether the reported memory is accurate or not. But sometimes we might be able to gather evidence that would help us evaluate such a claim in a rational way. Give some examples of ways in which we might be able to do this.
7. At the beginning of this course you took a "pretest" to see what you thought about various puzzles, pieces of reasoning, and the like. Why was it important to take this test *before* you studied the things in this course?

---

possible implications for false memories see Michael S. Gazzaniga, "The Split Brain Revisited," *Scientific American*, July 1998, 50–55. The case that the techniques discussed for improving eyewitness recall in fact improve it by 50% is made in Fisher and Geiselman, 1989 . . . Further references to be supplied.

8. Give some examples of the collective memory of people in the United States. Clearly some of the issues involving memory distortion in individuals are different from those in distortions in collective memory. But can you think of any similarities?
9. Think of a case where you thought that you remembered some event but later learned that the event had occurred rather differently from the way that you first thought. Write a paragraph describing this situation; include a discussion of things that might have led to the mistaken memories.
10. An eyewitness is (presented as) an authority or expert (on what they saw or heard), so the general questions we should ask about any authority are relevant. Imagine that you are on a jury and give some examples of the sorts of questions you would ask in trying to assess the accuracy of an eyewitness.
11. Explain what the *source amnesia* (i.e., misattribution of source) is, and give an example. Then discuss how it could be a danger in some concrete situation.

#### Answers to Selected Exercises

7. Most people are surprised when they are told, soon after they took the pretest, that certain answers are wrong (over 90% of the class usually gives the wrong answers to a few of the question). But if someone doesn't take the pretest, or takes it *after* learning about the relevant concepts, hindsight bias encourages them to think that they knew the correct answer all along and that they would have given the right answer on the pretest.

## 8.14 Appendix: Different Memory Systems

There is increasing evidence that memory involves a variety of subsystems. These systems serve different functions, work in different ways, and are located in different parts of the brain. For example working memory and long-term memory seem to take place in different parts of the brain.

Here is an intriguing example that bears on some of the topics in this chapter. The human brain consists of a left hemisphere and a right hemisphere. For the most part the left hemisphere is associated with the right half of the body, so it controls the right arm, and the right hemisphere is associated with the left half, so it controls the left arm. The input to the left half of each retina, which comes from the right side of the visual field, goes to the left hemisphere and the input to the right half of each retina, which comes from the left side of the visual field, goes to the right hemisphere.

Don't worry about the details here. The key points are that the left hemisphere controls the right arm and the right hemisphere controls the left arm. Furthermore, the left hemisphere "sees" the right visual field and the right hemisphere "sees" the left visual field. Language abilities are located primarily in the left hemisphere.

There are several bundles of nerves that connect the two hemispheres, and this allows them to communicate. But in some people these hemispheres have been severed (often in an attempt to prevent serious epileptic seizures that the person had suffered). In such cases neither hemisphere knows what is going on in the other one. The left hemisphere handles language so if just the right hemisphere sees something, a person with a split brain will not be able to *say* what it is, though he would be able to point to a picture of it with his left hand.

In a series of studies split-brain subjects were shown two pictures side by side in locations where the picture on the left is seen only by the right hemisphere and the picture on the right is seen only by the left hemisphere. If they are then shown a series of pictures, two of which related to the pictures they saw, each hemisphere would recognize the correct picture and direct the hand they control to point to it.

For example, if the picture on the left contains a snowman and the picture on the right contains chicken leg, the left hand will point to a picture of a snow shovel and the right hand will point to a picture of a chicken. (This can be a little confusing, but the picture in Figure 8.1 on the facing page should make it easy to sort things out.)

The person is then asked *why* the left hand is pointing to a snow shovel. Since language processing takes place primarily in the left hemisphere, the experimenter is asking the left hemisphere why the left hand is pointing to the shovel. Since the right hemisphere is responsible for pointing at the shovel, the left hemisphere doesn't know. Nevertheless, the left hemisphere will instantaneously make up






Figure 8.1: Where do False Memories Come From?

some reason, quite unconnected with the real reason, why the left hand is pointing to the shovel.

This suggests that false memories are created in the left hemisphere of the brain. This seems to be the part of the brain that searches for explanations, and it will often make them up if it can't find the real reasons. Further studies show that the left prefrontal regions of normal subjects are involved when people access false memories. This suggests that eventually we may be able to hook people up to a machine, a truth detector, to determine whether an apparent memory is accurate or false. This could have its uses, but the potential for invasion of privacy is appalling.

Research on these matters is all quite recent, and we still have much to learn. But there are two important implications for critical reasoning that are independent of the details.

1. Much of our thought and reasoning is not easily accessible to consciousness, and some isn't accessible at all. Mistakes in such reasoning may be harder to track down and much harder to correct than mistakes in conscious thought.

2. Memory consists of a number of distinct subsystems that serve different purposes and work in different ways. Indeed, the brain—and mind—in general seem to consist of a number of separate cognitive modules that perform quite different functions. In some respects your mind works more like a committee than like a single, unified agent.




Figure 8.2: The Line-up

## Chapter 9

# Emotions and Reasoning

**Overview:** Human reasoning never occurs in a vacuum; it is done by real people with needs and desires and feelings and emotions and moods. Many emotions are extremely valuable, and our goal should not be to set them aside every time we engage in reasoning. But intense emotions like anger and fear can cloud our judgment, and others people can exploit this to derail discussion or, worse, to manipulate us without our even realizing it. Moreover, some emotions actually provide an *incentive* to think badly in order to avoid unpleasant facts about ourselves or the world. In this chapter we will study some of the ways that emotions can impede clear thinking; we will then examine strategies to minimize their damage.

### Contents

---

<b>9.1</b>	<b>The Pervasiveness of Emotions</b>	<b>158</b>
9.1.1	Emotions and Information	158
<b>9.2</b>	<b>Stress</b>	<b>159</b>
<b>9.3</b>	<b>Legitimate Appeals to Emotion</b>	<b>160</b>
<b>9.4</b>	<b>Illegitimate Appeals to Emotion</b>	<b>161</b>
9.4.1	Pity	161
9.4.2	Fear	161
9.4.3	Anger	162
<b>9.5</b>	<b>Self-serving Biases</b>	<b>163</b>
<b>9.6</b>	<b>Chapter Exercises</b>	<b>165</b>

---

## 9.1 The Pervasiveness of Emotions

Emotions often *are* relevant to reasoning

Good thinking involves reasoning, not rationalization. It is based on what we have *good reasons* to think is true, not on what we would *like* to have be true. This shouldn't lead us to abandon emotions when we are thinking about things—we couldn't do so even if we tried. Emotions are a very central and important part of being human, and we are not like Mr. Spock on *Star Trek*, who is largely unaffected by feelings, nor would we want to be. Furthermore, emotions play a central role in *motivating* our actions. If we love someone that will lead us to treat them in certain ways; if we hate them we will treat them quite differently. If we lacked emotions, we wouldn't be motivated to do much of anything.

Emotions are a very mixed bag; they include joy, love, compassion, sympathy, pride, grief, sorrow, anger, fear, jealousy, envy and hatred. Although it is sometimes useful to speak as though emotions were one thing and thought was another, there really isn't a very clear line between the two. Emotions are not simply nonrational states that just happen to us. They can be more or less rational, more or less supported by the evidence, and they are susceptible to rational evaluation.

It makes sense, for example, to be angry in some situations but not in others. If I have evidence that Wilbur has punched Sam simply because he likes to hurt people, it makes sense to be angry with Wilbur. But if I'm angry with Wilbur just because I don't like his looks, anger doesn't make sense. In some cases jealousy has a basis in fact—a person really has taken up with someone else; in other cases someone dreams things up simply because they are insecure. Jealousy may not be a good emotion in either case, but in the first case it makes more sense that it does than in the second.

It is entirely reasonable to let our emotions and needs play a role in our plans and decisions; the fact that you love your child or take pride in your work gives you an excellent reason to take care of your little Wilbur and to do your job well. The fact that you are afraid of lung cancer gives you a good reason to stop smoking. The pity you feel for starving children is a good reason to donate money to charities that give them food. Indeed, without emotions you probably wouldn't care enough about anything to ever be able to act. But it's important to guard against the intrusion of emotions into places where they don't belong.

### 9.1.1 Emotions and Information

Emotions can affect all of the other things that we have examined earlier in this module.

**Perception** Our moods influence how attentive we are to things around us. People who are sad or depressed are often preoccupied with how they feel, and they focus less on the things around them. Emotions also affect our perceptual set, and so they color how we perceive things. For example, the people who watched the Princeton and Dartmouth football game (p. 79) probably saw it differently because of how they *felt* about their team. They identified with their university and its team, and most of people tend to see the groups they identify with as good.

If you find yourself home alone at night after watching a scary movie, shadows and sounds in the backyard or attic can assume new and sinister forms. Your fear and anxiety lead you to perceive things differently from the way you ordinarily do. Or if you have become jealous, a harmless and friendly conversation between your significant other and a friend may look like flirting.

**Testimony** If we like someone we may give too much weight to their testimony, and if we don't like them we may give too little. If we feel especially threatened or frightened by the world around us, it may even be tempting to give too much weight to the claims of demagogues or others who see conspiracies everywhere they turn. We *want* an easy answer to our problems, and that's just what they offer.

**Memory** We are constantly elaborating the information stored in memory. As we have seen, our emotions and moods can affect the ways we fill in details. For example our memories are sometimes selective and distorted in ways that protect our own self-image or self-esteem; we sometimes remember things the way we would like for them to be, rather than the way that they actually were.

**Fallacies and Biases** In the next chapter we will study several fallacies, ways in which reasoning can go wrong. Our emotions often give rise to fallacious reasoning. In still later chapters, we will see how emotions often lead to various self-serving biases in our thinking.

## 9.2 Stress

We all know that intense emotions like anger and fear can cause us to make all sorts of mistakes, so it isn't surprising that they can lead to flawed reasoning. Other emotions, like jealousy or grief, can also cloud our thinking. This is just common sense, but stress can pose a less-obvious, more long-term, danger.

**Stress** is an adverse reaction to the perception of a situation as harmful or threatening. It may involve physiological changes (e.g., tension headaches, sleep disturbances, trembling), and behavioral changes (e.g., inability to concentrate).

It can also impair the ability to think quickly or clearly. In extreme cases like panic people often suffer dramatic lapses in judgment. But even when stress is less severe it can lead to slips in reasoning, and it is never a good idea to make important decisions when under severe stress.

We should also be cautious in evaluating the claims people make when there are under a lot of stress. In the previous chapter we saw that eyewitnesses to crimes are less reliable than people commonly suppose. There are various reasons for this, but one seems to be that witnessing a crime, especially a violent one, produces stress. This in turn affects the witness's perception and memory and reasoning.

### Stress Management

It is not uncommon for college students who are away from home for the first time and facing many new challenges to have a problem with stress. There are various things one can do to **manage stress** better, including exercise, relaxation techniques, discussing one's problem with friends.

Most of these things take some time and effort, so it is easy to turn instead to people who promise a quick fix. In evaluating their claims we should ask the questions (listed in Chapter 5) that are always appropriate for evaluating self-styled experts. But help needn't be expensive. If you think that stress is a problem for you, you should consult a trained professional—for example someone at Goddard Health Center—who has experience dealing with it.

## 9.3 Legitimate Appeals to Emotion

Some emotions are typically positive and some are typically undesirable (happiness is usually good, but it's not healthy if you feel happy when you watch other people suffer). But most emotions are neither intrinsically good nor intrinsically bad. An emotion may be appropriate in one situation but not in another.

For example, anger is appropriate if we learn that someone has badly abused a subordinate, but it isn't appropriate if someone was three minutes late for an a meeting because they stopped to help the victim of a traffic accident. Weeks of grief are appropriate when a loved one dies, but not when your neighbor's pet gerbil does. So it shouldn't be surprising that there is nothing intrinsically wrong with people appealing to our emotions.

The person on the phone may appeal to my compassion and sympathy to get me to canvass my block for the March of Dimes. A coach appeals to her players' pride to spur them to play harder. A civil rights worker appeals to our sense of justice and fair play to induce us to behave differently toward members of a disadvantaged

group. The problems begins when someone appeals to emotions that aren't relevant in the situation, and it is especially serious when they manipulate our emotions without our realizing it.

## 9.4 Illegitimate Appeals: The Exploitation of Emotion

If we can't see a way to support our own view (or to refute someone else's view) using good arguments, it may be tempting to try to arouse emotions in the person we are arguing with on in third parties who are listening to our conversation. This diverts attention from the real issues so that people won't notice how weak our own case is.

Such diversion is particularly effective if the attack triggers intense emotions like anger or fear, because when we are angry or anxious it is harder to remain focused on the real issues and to think about them clearly. People may try to capitalize on our guilt or jealousy or envy or greed, but here we will focus on three of the most dangerous appeals, those to pity, fear, and anger.

### 9.4.1 Pity

Sometimes people appeal to our sense of pity. "It's true that I didn't do the homework for this course, but I've had a really bad semester, so can't you raise my *F* to a *D*?" "My client, the defendant, had a terrible childhood; you won't be able to hear about it without crying."

Appeals to pity and mercy are often legitimate, but they become problematic if we let our gut reaction dictate our response. We should instead evaluate the case on it's own merits. Perhaps the defendant did have a terrible childhood. But we have to stop and think about what this should mean, rather being moved solely by feelings of compassion or pity. In particular, the childhood is *not* relevant to the question whether the defendant is guilty of the crime. But it may be relevant in trying to decide what punishment is fair.

### 9.4.2 Fear

Fear affects our thinking, so if someone can arouse our fear they can influence what we think and, through that, how we behave. An appeal to fear is also the basis of the use of the scare tactic in advertising and other forms of persuasion. If someone can frighten us, that is a good way to make us draw a hasty conclusion without carefully evaluating the facts.

Appeals to emotions often involve exaggerations of various sorts, and an especially popular version of this is the **scare tactic**. The scare tactic aims to bypass

reason and manipulate us directly through our emotions. It plays on our fears, trying to convince us that we are in danger that can only be averted if we do what the other person suggests. It is common in advertising, including political advertisements.

1. We risk being social outcasts if we don't use a certain deodorant or mouth wash.
2. Life insurance commercials and tire commercials are especially adept at exploiting our anxieties and fears.
3. In politics negative campaigning is often combined with the scare tactic by alleging that some terrible thing will happen if a candidate's opponent is elected.
4. Demagogues try to exploit common fears and popular prejudices to entice us to support them; very often this involves placing the blame for our problems on others (e.g., members of another race or nationality).
5. One especially popular method in the age of the sound-bite is to use words that trigger emotions like anger and hatred and fear. Of course different words set off different people, but 'communist', 'atheist', 'bleeding-heart liberal', and 'redneck bigot' will be triggers for many.

### 9.4.3 Anger

One of the surest ways to derail an argument you are losing is to make the other person angry. They will then be more likely to lose sight of the real issues, and the fact that your case is weak will be forgotten once everyone has descended to accusations and name-calling.

For example, although debate over abortion is often conducted in a way that stays focused on the real issues, attacks on one's opponent are common here. Those who believe that abortion is permissible under some circumstances may be vilified as anti-life, cruel and heartless people, even murderers. Opponents of abortion may be said to be authoritarians who want to dictate how other people should live, as people who only too happy to trample all over a woman's right to decide what she does with her own body.

None of this means that emotions should be set aside when discussing abortion. The fact that we have the feelings about infants and human life that we do *are* relevant. But if open-minded discussion and mutual understanding are the goals, we have to discuss these things in a way that doesn't deteriorate into a shouting match.

When the prosecutor shows the jury photographs of the mutilated victim of grisly murder, he is appealing to emotions like anger and revulsion and shock. The



jury already knew the victim was murdered, but seeing the pictures arouses deeper feelings that mere descriptions of the murder scene ever could.

## 9.5 Self-serving Biases

Although there will always be people ready to exploit our emotions to further their ends, emotions and needs can lead us to reason badly without any help from anybody else. They can lead us to fool ourselves in order to avoid unpleasant facts about ourselves or the world. We cannot be effective thinkers if we won't face obvious facts or if we seriously distort them. We will study some of the mechanisms of self-deception at great length later on, but we should take a brief look at some of them now.

**Wishful Thinking** We engage in *wishful thinking* when we disregard the evidence and allow our desire that something be true to convince us that it really *is* true. True believers (p. 14) in a cause are especially prone to wishful thinking, but we are all susceptible, and in its more minor forms it is common.

The human tendency to wishful thinking is one reason why claims by pseudoscientists, advertisers, and others are accepted even when there is little evidence in their favor. There are many examples of this, and you can probably think of some from your own experience. For example, smokers find evidence that smoking is harmful to be weaker than nonsmokers do. People often greatly overestimate their chances of winning at games of chance or of winning a lottery (we will see later that the chances of winning a large state lottery are almost infinitesimally small).

**Defense mechanisms** Defense mechanisms are things we do, typically unconsciously, to keep from recognizing our actions, motives, or traits that might damage our self-esteem or heighten anxiety. Most defense mechanisms involve self-deception.

**Rationalization** Rationalization is a defense mechanism in which a person fabricates “reasons” after the fact to justify actions that were really done for other, less acceptable, reasons. We are all familiar with cases where people (probably even ourselves, if we think back on it) come up with a good “reason” for cheating on an exam or a diet, failing to do their homework, continuing to smoke despite their resolution to quit, or lying to a friend. Few of us like to view ourselves as dishonest, so if we do cheat a customer or lie on our tax return we are likely to rationalize it: everybody does it, they had it coming, they would have cheated me if they'd had half a chance, I really needed the money, and I'll never do it again.

**Repression** In the previous two chapters we studied the tricks memory can play. One of the easiest ways to avoid having to think about something is to simply forget it. It is unclear how frequent repression is. As we noted in the previous chapter, some childhood events (e.g., molestation) may be so traumatic and disturbing that people repress them. In recent years claims about repressed memories of childhood abuse have attracted a good deal of attention, but there is also some evidence that this occurs less often than people think, and some scientists think that many such reports are really cases of “false-memory syndrome.”

The important points here are that repression does seem to occur sometimes, but it is an empirical question how often it does. We can’t answer such questions by taking a vote or by going with our gut feelings. We can only answer them by a careful consideration of the relevant evidence.

**Denial** Denial is a refusal to acknowledge the existence or actual cause of some unpleasant feature of ourselves or the world. Here unacceptable impulses and disagreeable ideas are not perceived or allowed into full awareness. Denial is a defense mechanism that is far from rare (therapists like to say that Denial is not just a river in Egypt). For example, it is common for those with serious drug or alcohol problems to deny (even to themselves) that they really have a problem (“I could quit any time I wanted to”). Often those close to the person engage in denial too.

**Self-deception** Self-deception occurs when someone fools himself into believing something that is not true. For example, many people have unrealistically high opinions of themselves. People often engage in self-deception to boost their ego or enhance their self-esteem, but they may do so for other reasons as well. For example, a mother may be unable to believe that her son has a drug problem, even though she has found syringes in his room several times.

Wishful thinking, rationalization, and denial shade off into one another, and we won’t worry about making fine distinctions among them. It is an empirical question just how widespread they are, but there is good evidence that they are common. What is absolutely clear is that they pose problems for clear and accurate thought. They all lead us to ignore what is really going on, which means that we can’t reason about it clearly.

**Lake Wobegon effect:** well over half of us think we are above average in various ways

**The Lake Wobegon Effect** A large majority of adults in this country think that they are above average in a variety of ways, and only a very small percentage think that they are below average. For example a survey of a million high-school seniors found that 70% rated themselves above average in leadership skills while only 2% felt they were below average. And *all* of them thought that they were above average

in their ability to get along with others. Most people also think of themselves as above average in intelligence, fairness, job performance, and so on through a wide range of positive attributes. They also think they have a better than average chance of having a good job or a marriage that doesn't end in divorce.

This finding has been called the Lake Wobegon effect after the fictional town Lake Wobegon in Garrison Keillor's "Prairie Home Companion," a place where "the women are strong, the men are good-looking, and all the children are above average."

**Self-serving Biases** All of these things can promote self-serving biases. We will see many examples in later chapters, so one example will suffice now. People have a strong tendency to attribute their successes to their own positive features (good character, hard work, perseverance) while attributing their failures to external conditions beyond their control (bad luck, other people didn't do their share of the work). "I did well on the first exam because I'm bright and I studied really hard." "I did poorly on the second exam because I felt sort of sick, and besides the exam wasn't fair." As we will see, we aren't usually so charitable with others. "I was late to work because the traffic was really bad." "Sam was late to work because he just can't get it together to organize his time."

Emotions are an important part of life, in many ways *the* most important part. But as we have seen in this chapter, they can also cloud our reasoning in ways that are harmful to others, and to ourselves.<sup>1</sup>

## 9.6 Chapter Exercises

1. Explain the role that emotions played in some of the arguments used by people who supported the view that the six-year old Cuban, Elian Gonzalez, should be kept in the U.S. Then explain the role that emotions played in some of the arguments used by people who supported the view that Elian Gonzalez should be returned to his father.
2. Should a prosecutor be allowed to show the jury grisly photographs of a murder victim? What reasons can you think of for not allowing the pictures to be displayed. What reasons can you think of for allowing it? Would it make a difference whether

---

<sup>1</sup>You can find a very accessible discussion of self-deception and related matters in Daniel Goleman's *Vital Lies, Simple Truths: the Psychology of Self-Deception* Simon & Schuster, 1985. A good discussion of the Lake Wobegon effect and references to empirical studies of it can be found in Thomas Gilovich's *How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life*, The Free Press, 1991. We will consider people's attributions of their success and failures in more detail when we examine the fundamental attribution error. Further references will be supplied.

the pictures were shown during the first part of the trial (before the jury has found the defendant innocent or guilty) of in the sentencing phase (after they have found her guilty and are trying to decide on the appropriate punishment)?

3. Give an example of someone else's engaging in self-deception or wishful thinking. What do you think leads them to do this and how might they avoid it?
4. Give an example of an appeal to pity. Should it move us? Under what conditions are such appeals legitimate? Under what conditions do they seem inappropriate?
5. Discuss some ways in which wishful thinking has affected your own thought (or those of others). How might you (or they) have avoided its unhealthy effects. Can you think of cases where wishful thinking might lead to good outcomes?
6. Analyze each of the following dialogues.

*Edna:* So, how'd the logic class go?

*Wilbur:* It really sucked.

*Edna:* What grade did you get?

*Wilbur:* I flunked. But it wasn't my fault. The teacher was a complete loser. Anybody who passed that course would have to be a real idiot.

*Edna:* I'm sorry to have to put it like this, but since you just keep pushing, you don't really leave me any choice. I just don't *want* to go out with you. I'm sorry.

*Wilbur:* Well, I didn't want to go out with you either. I just asked you out because I felt sorry for *you*.

*Edna:* Isn't that your eleventh beer this evening?

*Wilbur:* What's it to you? It's been a lousy semester, what with my pet hamster Emmy Lou dying and that disgusting logic class. So I deserve to unwind a little. Anyway, I could quit drinking any time I wanted—if I wanted.

*Logic Teacher:* Please put your homework in the "In" folder.

*Wilbur:* I pulled a real late nighter and finished all of the homework. But after I typed it all up and saved it to disk, my computer crashed, and I lost it all.

*Logic Teacher:* That's the third time this semester.

*Wilbur:* I know. It's like the computer's out to get me.

*Edna:* How did the job interview go?

*Wilbur:* It went really well. I have a really good feeling about it.

*Edna:* But didn't you feel the same way after those other interviews you had, the ones where they never called back? How many was it anyway, eighteen?

*Wilbur:* Twenty. But I really really feel good about this one. I just know I'll get it.

7. In what sorts of situations or circumstances is it reasonable to let our emotions influence us; in what ones is it not such a good idea? Give some examples of each and defend your choices.
8. Think of an instance in your own life where you later felt that you may have used a defense mechanism, perhaps to boost your ego. What might have led to the self-deception. (You will not be asked to hand this in, but it's worth a few moment's thought.)
9. Discuss some ways in which strong feelings of guilt might impair clear thinking.
10. Write a dialogue that illustrates the bad effects of self-deception on reasoning. Then write a second dialogue that illustrates wishful thinking, and a fourth that illustrates denial.



## **Part IV**

# **Relevance, Irrelevance, and Fallacies**





## Part IV. Relevance, Irrelevance, and Fallacies

In reasoning and argumentation it is important to stay focused on the topic at issue. This means giving reasons or evidence that *bears* on the topic, that is *relevant* to it. This sounds easy, but a great deal of bad reasoning occurs because we don't stay focused. In chapter 10 we study relevance and several common ways in which reasoning goes awry when we use premises or evidence that aren't relevant.

Bad reasoning is said to be *fallacious*. Some fallacies are simple to spot, but others give the appearance of being good arguments, and it is easy to be taken in by them. In chapter 11 we learn about three fallacies: begging the question, the either/or fallacy and, more briefly, the fallacy of the line. We will also note problems involving inconsistency.



## Chapter 10

# Relevance, Irrelevance, and Reasoning

**Overview** In reasoning and argumentation it is important to stay focused on the topic at issue. This means giving reasons or evidence that *bears* on the topic, that is *relevant* to it. In this chapter we study relevance and several common fallacies of relevance; these are common ways in which reasoning goes awry when we use premises or evidence that are not relevant to our conclusion.

### Contents

---

<b>10.1 Relevance</b> . . . . .	<b>174</b>
<b>10.2 Fallacy of Irrelevant Reasons</b> . . . . .	<b>178</b>
<b>10.3 Arguments Against the Person</b> . . . . .	<b>180</b>
<b>10.4 The Strawman Fallacy</b> . . . . .	<b>183</b>
10.4.1 Safeguards . . . . .	185
<b>10.5 Appeal to Ignorance</b> . . . . .	<b>187</b>
10.5.1 Burden of Proof . . . . .	188
<b>10.6 Suppressed (or Neglected) Evidence</b> . . . . .	<b>191</b>
<b>10.7 Chapter Exercises</b> . . . . .	<b>192</b>

---

## 10.1 Relevance

In reasoning and argumentation it is important to stay focused on the topic at issue. This means giving reasons or pieces of evidence that *bear* on the topic, that *pertain* to it, that are *relevant* to it. This sounds easy, but a great deal of bad reasoning occurs because we don't stay focused on the issues.

Relevance is important in all communication, even when we are not constructing arguments or trying to persuade other people. In a normal conversation each person typically says things that are relevant to the general topic and to what the other person has been saying. Slight departures from relevance are all right, but someone who keeps bringing in irrelevant points, things from out of left field, is difficult to talk to. If they act this way a lot of the time, they even earn unflattering labels like "space-case."

Relevance is a *relationship* between sentences

A statement is not relevant all by itself, in isolation from anything else. Relevance instead involves a *relationship* between one statement and another. So a premise can be relevant to one conclusion, but completely irrelevant to others. It is *irrelevant* if it simply *doesn't bear* on the truth or falsity of the conclusion, if it is *independent* of it, if it *does not affect it one way or the other*.

### Examples of Relevance

1. The premise that witnesses claim to have seen Timothy McVeigh rent the Rider truck used in the bombing of the Alfred P. Murrah Federal Building in Oklahoma City is relevant to the conclusion that he is guilty.
2. The premise that Wilbur has failed his first two exams in Chemistry 1113 is relevant to the conclusion that he will fail the course.
3. The premise that the death penalty deters murder is relevant to the claim that we should retain capital punishment for murder.
4. The premise that the death penalty does not deter murder is relevant to the claim that we should retain capital punishment for murder.

### Examples of *Irrelevance*

1. The fact that over a hundred and sixty people were killed in the Oklahoma City bombing is *not* relevant to the claim that Timothy McVeigh is guilty (though in the sentencing phase of the trial, after he was convicted, it may have been relevant to the claim that he should receive the death penalty).

2. When a reporter asks a politician a tough question, the politician often gives a long response that really doesn't answer the question at all. Here the politician acts as though what she is saying supports her position, but it may be completely irrelevant to it. This happens when she puts a "spin" on things that shifts the focus from the real issue to something else.
3. Many advertisements use endorsement from celebrities. Often the fact that a famous person endorses a something is altogether irrelevant to the claim that it's a good product. For example, the fact that Michael Jordan plugs a certain cologne isn't likely to be relevant to the conclusion that it's a good cologne.
4. In everyday conversations people are often at "cross purposes," they "talk past each other." This can occur when they *think* they are discussing the same claim or issue, but in fact the two of them are actually concerned with somewhat different issues. In such cases, the things each person says in support of their own views may seem irrelevant to the other person.

### Relevance vs. Other Concepts

#### Relevance is not the same thing as truth

- A premise can be true, but irrelevant to a given conclusion.

Example: it is *true* that Timothy McVeigh has a certain astrological sign, but this is irrelevant to the claim that he is guilty of the bombing.

#### Irrelevance is not the same thing as falsity

- A premise can be false, but still be relevant to a given conclusion. To say that it is relevant is to say that *if* it were true, it would make the conclusion more (or less) likely to be true.

Example: The claim that McVeigh wrote a letter to the FBI saying he would bomb the Murrah Building is false. But it is relevant to the claim that he is guilty, since if he had written such a letter, it would make it more likely that he was guilty.

#### Relevance is not the same thing as importance

- An *important* claim can be *irrelevant* to a given conclusion.

Example: It is a very important fact that many people were murdered in the bombing, but this is irrelevant to the conclusion that McVeigh committed the murders.

- An *unimportant* claim can be *relevant* to a given conclusion.

Example: The fact that there are over ten thousand blades of grass in Wilbur's lawn isn't very important to anyone, but it is highly relevant to the claim that there are over nine thousand blades of grass in his lawn.

### Relevance is not the same thing as conclusive support

Relevance comes in degrees. It is not an all-or-none matter. Some premises are highly relevant to a given conclusion, others are somewhat relevant, and yet others are completely irrelevant. So to say that a premise is relevant to a conclusion is *not* to say that it provides conclusive support for the conclusion.

**Relevance can be either positive or negative** Any claim that provides evidence for, or against, some other claim is relevant to it. It has positive relevance if it supports it or counts in favor of it. It has negative relevance if it makes it less likely or counts against.

- the claim that John's finger prints are on the murder weapon is relevant to the conclusion that he committed the crime. It makes it more likely, and so has *positive relevance* for this conclusion.
- The claim that John was seen in another state at the time of the murder is also relevant to the conclusion that he committed the crime. It makes it less likely, and so has *negative relevance* for this conclusion.

If two claims are irrelevant to each other they are sometimes said to be *independent* of each other. The truth-value of one has no effect or influence or bearing on the truth value of the other. Knowing that one is true (or false) tells you nothing whatsoever about whether the other is true (or false). Irrelevance is a two-way street: if one thing is irrelevant to a second, the second is also irrelevant to the first.

- Example 1: If you are flipping a fair coin, the chances that it will land heads on any given flip is  $1/2$ . The outcomes of successive flips are independent of each other, so the outcome on the previous flip is irrelevant to what you'll get on the next flip.
- Example 2: If you and your spouse are going to have a child, the chances that it will be a girl are very nearly  $1/2$ . The outcomes of successive births are independent of each other, so the sexes of your previous children are irrelevant to the sex of your next child

**Exercises**

1. Write a sentence or two explaining why each of the following claims is either relevant or irrelevant (as the case may be) to the claim that we should abolish the death penalty in Oklahoma.

1. innocent people have sometimes been executed
2. statistics show that it deters (decreases) murder
3. some people really enjoy watching executions
4. statistics show that it does not deter murder
5. it makes the warden sick to his stomach

2. Say whether the members of the following pairs are positively relevant, negatively relevant, or simply irrelevant to each other, and say why.

1. President Clinton resigns. The Vice President becomes President.
2. Jessie “The Body” Ventura is elected President in the next election. I roll a 3 on the first roll of a die.
3. Jessie “The Body” Ventura is elected President in the next election. The economy goes south.
4. I get a head on the next flip on a coin. I get a head on the flip after that.
5. I pass all of the exams in this course. I pass the course itself.
6. I miss a lot of classes. I pass the course.
7. There are many valuable things about sports. Names of sports teams like “Redskins” aren’t demeaning to Native Americans.

3. Give two premises that are relevant (in the sense discussed in class and in the text) to the following conclusions. Then give two premises that are irrelevant.

- a. We should not send ground troops into the Balkans
- b. President Clinton should have been impeached
- c. I’ll probably not get an “A” in my chemistry course.

**Answers to Selected Exercises**

1. Write a sentence or two explaining why each of the following claims is either relevant or irrelevant to the claim that we should abolish the death penalty in Oklahoma.

1. innocent people have sometimes been executed

- Relevant (various analyses possible, but you need to defend your answer)
2. statistics show that it deters murder
    - Relevant (various analyses possible, but some analysis needed). Note that relevance can be either positive (supporting a view) or negative (weakening the case for it).
  3. some people really enjoy watching executions
    - Irrelevant (various analyses possible, but some analysis needed)
  4. statistics show that it does not deter murder
    - Relevant (various analyses possible, but some analysis needed). Note that relevance can be either positive (supporting a view) or negative (weakening the case for it).

## 10.2 Fallacy of Irrelevant Reasons

**Fallacies:** common and tempting ways of reasoning badly

If the premises of an argument are irrelevant to the conclusion, then the argument is flawed. The premises may well be true, important, and perhaps even relevant to other conclusions we care about. But if they aren't relevant to the conclusion we are thinking about, then the argument is bad. Bad reasoning is said to be *fallacious*, and some bad patterns of reasoning—fallacies—occur so frequently that it is useful to give them names. This helps us spot them and to avoid them in our own thinking.

**Fallacy of irrelevant reason:** using an irrelevant premise to support a claim

We commit the **fallacy of irrelevant reason** (or irrelevant premise) if we offer a premise to support a conclusion when the premise is *irrelevant* to the conclusion. Relevance is *not sufficient* for premises to support a conclusion, but it *is necessary*. The fallacy is also known by its Latin name, *non sequitur* (“it doesn't follow”). So when someone draws a conclusion from irrelevant premises, we say that it's a *non sequitur*.

The name *irrelevant reason* is a sort of catch-all label. All of the fallacies that we will study in this chapter have premises that are irrelevant to the conclusion. But if we have a more specific name for a fallacy (and we won't until we get further into this chapter), we will use the more specific label.

### Motivations behind the fallacy

If we can't supply relevant reasons to support a conclusion, it is tempting to bring in something that is irrelevant. This deflects attention from the fact that we don't have good reasons for our view. This is especially effective if the irrelevant reasons



have emotional impact, because these make it particularly easy to focus on things that are not relevant to the real issue.

The fallacy of irrelevant reasons is sometimes called the **red herring** fallacy. It gets its name from the fact that people who were fleeing from trackers with bloodhounds would sometimes wipe a dead animal across the path to throw the dogs off their trail.

Sometimes responses are so irrelevant that they really don't look like reasons at all, but they can still deflect attention from the real issue. Jokes, ridicule, sarcasm, flattery, insults, and so on can deflect attention from the point at issue. A joke can be particularly effective, since if you object that it isn't relevant, you can then be accused of lacking a sense of humor. But you can laugh at the joke and then return to the issue.

### Safeguards

1. Whether or not reasons are relevant to a conclusion depends on the conclusion and the way it is stated. So always stay focused on the conclusion.
2. Don't allow jokes, insults, or the like to deflect your attention from the issue. You can appreciate the joke, but then return to the point at issue.
3. Be sure that you and the person you are talking to really are considering the same claim rather than talking at cross purposes. If you are, try to explain your view to the other person *before* defending it. This doesn't guarantee the two of you will come to an agreement, but the discussion will be more productive.

### Exercises

1. Identify the fallacy (if any) in the first two passages.
  - O. J. Simpson is innocent. He had an alibi and he loved Nicole.
  - O.J. Simpson is innocent. He's a very famous person, and was a fantastic football player.
2. Which of the following are relevant to the conclusion that we should *not* have laws making handguns harder to get?
  - (a) The Bill of Rights says that we have a right to bear arms.

- (b) Many people have avoided serious injury because they had a gun and were able to frighten off an intruder.
  - (c) Many of the people who favor gun control are really just frightened by guns.
  - (d) Many children are accidentally killed each year by guns in their homes.
  - (e) It's a lot of fun to go out and try to shoot stop signs with a handgun.
3. Give several premises that are relevant to the following conclusion. Then give several that are irrelevant:

Abortion should be illegal (except when the mother's life is in danger).

### 10.3 Arguments Against the Person

**Argument against the person:** attacking a person rather than their argument

We commit the **fallacy of an argument against a person** whenever we launch an *irrelevant* attack on that person, rather than on her position or argument. The Latin name for this fallacy, **ad hominem**, is still in common use, so we will use it too. This is one type of the fallacy of irrelevant reason, since when we attack a person we shift our focus from issues that are relevant to the conclusion to another issue that is not relevant; in this case we shift our focus to the person we are attacking.

If we disagree with a position, or if an argument has a conclusion we reject, it is perfectly reasonable to try to show that the position is false or that the argument is flawed. But when we can't see a way to do this, it may be tempting to instead attack the person who holds the position or who gave the argument.

This diverts attention from the real issue, shifting the focus elsewhere, so that people won't notice how weak our own case is. And one of the most effective ways to shift the focus is to attack the other person in a way that triggers various emotions like anger, because when we are angry or incensed, it is difficult to remain focused on the real issue.

The simplest way attack a person is to simply throw terms of abuse at them. These range from 'fool' or 'idiot' to derogatory labels based on the person's race or nationality or gender or sexual orientation. We are all familiar with cases where discussion or debate degenerates into name calling.

For example, debate over affirmative action programs is often conducted in a way that stays focused on the real issues, but attacks on one's opponent are not uncommon here. Champions of affirmative action are sometimes accused of being bleeding heart liberals who really want to discriminate against white males, while opponents of affirmative action are sometimes accused of being rednecks or bigots who only want to hold on to a situation that benefits them at the expense of others.

One label has undergone an interesting reversal of fortune; several decades ago, many considered it a good thing to be a *liberal*. The term originally signified those who favored liberties and freedoms (for example freedom of religion). But in this century, partly as a result of Roosevelt's New Deal, many people came to see liberals as champions of a "tax and spend" approach to government. So when Dole accused Clinton of being a liberal in the 1996 Presidential debates, Clinton was quick to disavow the label (saying that it was a tired "golden oldy," and that "that dog won't hunt" anymore).

It is also possible to attack someone by pointing out that they are associated with a group we don't like. Such attempts to show guilt by association commit the *ad hominem* fallacy if they take the place of a reasonable examination of the other person's argument or views.

A more subtle version of the *ad hominem* fallacy occurs when we ignore someone else's argument for a given position and instead charge that they only favor the position because it is in their *self-interest* to do so. The following dialogue represents a typical instance:

*Burt:* Well, anyway, there you have my arguments for opposing gun control laws.

*Al:* Well, all those fancy statistics and detailed arguments sound good. But when you get right down to it, you really oppose gun control because you see guns, and you'd lose a bundle if any laws were passed that cut back on their sales.

Here Al has simply ignored Burt's arguments for his position and attacked Burt instead.

In some cases attacks on a person may be hard to resist. Suppose, for example, that someone gives us a good argument, based on lots of statistics, that we should wear seat belts. Later we learn that they always ride their motorcycle without a helmet. This does show some inconsistency, and perhaps even hypocrisy, in their behavior. But it doesn't show that their argument for wearing seat belts is bad.

Not all "attacks" on a person are irrelevant. If someone purports to be a good *source of information* about something, it is perfectly reasonable to expose them if they really aren't a good source.

Example 1: If someone purports to be a highly trained expert in some field (e.g., they claim to have a medical degree) when in fact they lack the training they claim to have, this is worth noting and it does damage their credibility.

Example 2: If a source has repeatedly been wrong, for example if a checkout-counter tabloid has frequently be wrong in its claims about

Hollywood stars, then it is not a good source of information. Here it is relevant to point out that the source has a poor track record, since that should affect our assessment of their current claims.

Example 3: If an eyewitness to a murder is testifying in court, it is reasonable to offer testimony to show that her eyesight is poor or her memory faulty or that he has a reason to lie.

Example 4: If a person or group has repeatedly shown biases or prejudice about certain issues or against other groups, it is unwise to trust them when they make further claims about those issues or groups. When it became clear during the O. J. Simpson trial that the policeman Mark Furhman had repeatedly used the word ‘nigger’, his testimony came into serious doubt.

We only commit the *ad hominem* fallacy if we ignore someone’s arguments or reasons and instead attack them. Of course life is short, so if someone is known to be biased or unreliable, that *does* justify spending our time doing better things than thinking about their argument. But it does *not* justify concluding that their *argument* is no good.

### Exercises

Say whether or not each of the following passages contains an *ad hominem* fallacy. If it does, explain how the fallacy is committed and say how an attack on the argument, rather than on the arguer might proceed.

1. You just argue that OU should adopt a pass/fail grading system so you won’t have a bunch of Ds.
2. Well, of course professors here at OU can cite all sorts of studies and give all sorts of arguments that they deserve a pay increase. After all, they are trained to do that stuff. But at the end of the day, they are just like the rest of us, looking out for old number one.
3. You know the Pope’s arguments against birth control. But, you know, I say if you don’t play the game, don’t try to make the rules (Dick Gregory).
4. The witness of the defense is hard to take seriously. He testifies as “expert witness” for the defense in over a hundred trials a year, and they always pay him big bucks to do it.
5. Who are you to tell me not to smoke a little dope. You knock off nine or ten beers a day “to relax.”

6. Look, I hear all your arguments that abortion is wrong. But you're a man, and you can't be expected to understand why a woman has to have a right to choose what to do with her own body.

## 10.4 The Strawman Fallacy

We commit the **Strawman Fallacy** when we distort or weaken someone's position or argument in an effort to discredit it. When this happens, we are not really countering the person's actual views, but are merely assailing a feeble version of them. We are said to be **attacking a strawman** (it might seem more accurate to say that we are attacking a straw argument, but the word 'strawman' is the traditional label for this fallacy).

**Strawman fallacy:**  
distorting someone else's position to make it easier to attack

This is another type of fallacy of irrelevant reason, since when we attack a weakened version of a view or argument, we shift attention from issues that are relevant to the conclusion, the real argument that is given for it, to other issues, the weakened caricature of the argument.

When we are confronted by a position that conflicts with our own, it is often tempting to characterize the position in the weakest or worst or least defensible light. We make our own view *look* strong by making the alternative look weak (rather than showing that our view *is* strong by building a solid case for it). By distorting the opposing position, we make it easier to answer, or even to dismiss, the view and those who subscribe to it. This saves us the trouble of having to think seriously about it, and spares us the possibility of having to acknowledge that we might be wrong. These are scarcely good things to do, but perhaps even worse, we show the other person a lack of respect by not taking them seriously; we don't like it when someone else does this to us, and others won't like it when we do it to them.

### Examples

Many campaign ads, especially "attack ads," go after strawmen. For example, someone who opposes the death penalty is likely to be accused of being soft on crime or of favoring the rights of criminals over the rights of victims. Someone who favors certain welfare programs may be accused of wanting to tax and spend in order to give people handouts and discourage their sense of responsibility, while someone who favors cutting back on the same programs may be accused of trying to help the rich at the expense of the poor. And a person who advocates decriminalization of drugs may be accused of favoring the use of drugs.

But politicians are not the only people who commit this fallacy; all of us do it at

one time or another. It's not difficult to find people who say things like "I'm sick of all those self-righteous people arguing that Oklahoma shouldn't have a state lottery. They just think that anything enjoyable is a sin." While there may be some killjoys who oppose the lottery on the grounds that it's fun, most of its critics oppose it for more serious and substantial reasons. For example, parents may spend money that their children badly need on lottery tickets.

Attacks on a person are similar to attacks on a strawman insofar as both ignore a person's actual argument or position. But they differ in an important way:

1. Someone commits the strawman fallacy when they ignore a person's actual argument and attack a weaker, distorted version of it.
2. Someone commits the *ad hominem* fallacy when they ignore a person's actual argument and attack the person instead.

Like the other fallacies, we may commit the strawman fallacy intentionally. But human nature being what it is, it is all too easy to commit the fallacy without really thinking about it. We do it, for example, if we thoughtlessly restate our spouses' views about something in a way that makes them seem less plausible or compelling than they really are.

### More Subtle Versions of the Strawman Fallacy

There are several special cases of the strawman fallacy that can be especially difficult to detect.

#### Taking Words out of Context

Taking someone's words out of context also allows us to quote them in a way that can make their position look weaker than it really is. We can leave out the qualifications and complexities of the view that would allow it to withstand the criticisms that we direct against the weaker version.

For example, a Senator who voted for a bill containing many provisions (including a small tax increase which, taken alone, she opposed) might be characterized as favoring higher taxes.

#### Treating an Extreme Case as Representative

It can be particularly effective to treat the views of an extreme member of a group as representative of those of the entire group, since this allows us to literally quote someone. We use their own exact words, in a way that appears to convict the entire group. The fact that we use their own words makes it seem more likely that we are

being fair. Since we don't like the view, perhaps we couldn't be expected to give a fair summary of it, but here we seem to have the view right from the horse's mouth.

For example, opponents of gun control laws can find people who would like to ban all guns, and they may quote their views as though they were representative of the views of all people who favor some restrictions on guns. At the same time, people on the other side of this issue sometimes quote members of the more extreme militia groups as though they were representative of all of those who think gun control laws can be overdone.

### **Criticizing an Early or Incomplete Version of a View**

Criticizing an early or sketchy version of a view, rather than considering it in its current, stronger form, also make it much more vulnerable to attack.

For example, Burt might attack the theory of evolution by quoting Darwin and showing that his views on some detail of the theory are now known to be wrong. But Darwin wrote well over a century ago, and the theory has undergone many refinements and improvements since his day. You may not like the theory, but if you want to show that its wrong, you have to consider the strongest version of it.

### **Criticizing a Deliberately Simplified Version of a View**

Sometimes people state their position in a simplified way in order to get their basic ideas across in a short time. If it is clear that someone is doing this, their opponents should go after the more complex version of their views.

#### **10.4.1 Safeguards**

Several things that can help us to spot, and to avoid committing, the strawman fallacy.

1. Be aware of the natural human tendency to characterize opposing views in a way that makes them easier to attack or dismiss.
2. Be fair. Try to find the strongest version of the view in question and consider it. Give the other person the benefit of the doubt. This will require you to think harder, but if you do, your own views and your reasons for holding them will be more secure.
3. Do not rely on the critics of a view to state the view fairly. They may do so, but you can't count on it. This is especially true when the point at issue is highly controversial or arouses intense emotions.

**Exercises**

In each of the following passages,

1. Determine whether it contains a strawman fallacy (is there an attack on a strawman?). If there is,
  2. Explain the way in which it commits the fallacy.
  3. Note ways in which the passage could be revised so as to be more fair to the view under consideration.
- 
1. The problem with people in the environmentalist movement is that they lack common sense. They think that protecting the environment for the spotted owl is more important than allowing people to make a living cutting timber in the owl's habitat.
  2. A bill that would allow school prayer would be very bad. It's supporters would like to have everyone involved in religion, and in fact in the Christian religion. This would violate the rights of those who aren't Christians.
  3. People who are against school prayer really want to get rid of all religion. At bottom they are atheists, or at least agnostics.
  4. Champions of campaign finance reform have a hopeless view. They seem to think that if there were limits on contributions to political candidates there would be no more corruption in politics and the poor could also afford to run for office. But we will always have corruption, and it will always be easier for the rich to get elected.

**Answers to Selected Exercises**

You were asked to determine whether each passage contains a strawman fallacy (is there an attack on a strawman?). If it did, explain the way in which it commits the fallacy. Finally, note ways in which the passage could be revised so as to be more fair to the view under consideration. You can think about the last part for yourself; here are answers to the first two parts

1 The problem with people in the environmentalist movement is that they hold that protecting the environment for the spotted owl is more important than allowing people to make a living cutting timber in the owl's habitat.



This passage attacks a strawman. A few environmentalist may think this, but most do not hold the extreme position that the owl is more important than human livelihood.

2. A bill that would allow school prayer would be very bad. It's supporters would like to have everyone involved in religion, and in fact in the Christian religion. This would violate the rights of those who aren't Christians.

It may be that this characterization of the position of proponents of school prayer accurately captures the position of a few of them. But most people who favor school prayer do not have any extreme view of this sort, so the passage attacks a strawman.

3. People who are against school prayer really want to get rid of all religion. At bottom they are atheists, or at least agnostics.

The problem here is similar to the problem in 2, but now the fallacy is being committed by people on the other side of the school-prayer issue.

## 10.5 Appeal to Ignorance

Every Halloween night, in the comic strip "Peanuts," Linus makes his yearly pilgrimage to a local pumpkin patch to await the Great Pumpkin's arrival. Many of his friends are skeptical (although Sally usually accompanies him), but Linus remains convinced.

Now suppose someone offered you \$50 to prove, right here on the spot, that the great pumpkin does not exist. Could you do it? Could you even come up with good evidence to show that the Great Pumpkin probably doesn't exist? I can't. But if you can't, does that mean that you should see the issue as an open question, that you should regard it as a 50/50 proposition? No. Linus is a kid, but suppose he had to enter the real world, grow up, and go off to college. What would you think about him if he still believed in the Great Pumpkin when he was 32? What would you think if you arrived at OU and found that your new roommate believed in the Great Pumpkin?

What is the moral of this story? Most of us cannot give strong evidence that the Great Pumpkin doesn't exist, but we would regard anyone who thinks that it



is a completely open question as much too gullible. Of course we don't encounter adults who believe in the Great Pumpkin. But we all encounter people who make some claim that seems implausible. Then, instead of building a positive case to support his claim, he suggests that since we can't show that it's wrong, it is probably true. We all have heard the refrain: "Well, you can't show I'm wrong ..."

### 10.5.1 Burden of Proof

When you make a claim that everyone agrees is true (e.g., that OU's main campus is in Norman), you don't need to build a case for it. When everybody already thinks something is so, you don't need to give reasons to show that it is. But if you make a surprising or controversial or implausible claim (e.g., that several OU students have been abducted from campus by aliens from Mars), then it is your *responsibility* to give reasons for your claim. And the more implausible the claim, the heavier your burden of proof. So the fact that you can't produce evidence that the Great Pumpkin does not exist gives you absolutely no reason to think that it really does.

When someone defends a view by pointing out that you can't show that it's false, they are committing the **fallacy of appeal to ignorance**. The fact that you are ignorant (don't know) of evidence that would show they are wrong does not mean they are right. This is a fallacy of irrelevance, since the fact that I cannot show that some claim is false is not relevant to showing that it is true.

When someone makes a surprising claim, then adds, "Well, you can't show that I'm wrong," they are unfairly shifting the burden of proof to you. We often are in no position to prove that their claim is false. For example, if someone makes a claim about aliens from outer space infiltrating critical reasoning courses at OU, I cannot prove that there haven't been any. How could I? But the claim is implausible, and *until* someone gives us reasons to believe it, it's reasonable to believe that it is *false*.

This bears repeating. The reasonable attitude here is not complete open-mindedness. It is *not* sensible to conclude that it's a fifty/fifty proposition that creatures from outer space are stalking our campus. Until we are given some reason to believe this claim, it is much more reasonable suppose that it is false.

There are many cases like this. You are not now in any position to show that the Great Pumpkin doesn't exist. But if you went around thinking that it was a fifty/fifty proposition that there was a great pumpkin, people would have serious doubts about you (and well they should). Of course most of us aren't worried about the Great Pumpkin. The fallacy is worth studying because there are many other, less obvious, cases of the same sort. In short: absence of evidence for *X* is false is not evidence that *X* is true. The fact that I can not cite conclusive evidence for my view that there is not a Great Pumpkin is not evidence that my view is false.

**Burden of proof:** when someone makes a surprising claim, it's *their* job to defend it

Absence of evidence that *X* is false  $\neq$  evidence that *X* is true

Note that an appeal to ignorance does not involve saying that someone else is ignorant or misinformed or just plain dumb. The word ‘ignorance’ has a special meaning here. Someone commits the fallacy of an appeal to ignorance when they suggest that the fact that they haven’t been shown to be wrong is somehow evidence that they are right.

### Positive vs. Negative Claims

Let’s call a claim that there are *Xs* a *positive existence claim* and a claim that there are not any *Xs* a *negative existence claim*. To show that a positive existence claim is true, it suffices to point to an example of *X*. If a biologist claims that she has discovered some new, unsuspected strain of virus, she can prove her case by producing a sample of it and allowing other scientists to test it.

But it can be very difficult to prove that a negative existence claim is false, particularly if it says that there are no *Xs* anywhere at all. For example, you cannot really look everywhere to determine that there are no *Xs*; you can not look everywhere and then report that there wasn’t a Great Pumpkin anywhere you looked. Nevertheless, the claim that there is one is implausible; no credible witnesses have seen it, and science gives us no reason for believing in it.

### Open-mindedness

Open-mindedness is a very great virtue, but it does not require us to seriously entertain any claim that comes down the pike. It does require us to remain ready to reevaluate any of our beliefs if new evidence or arguments come along, and being willing to change our beliefs if the evidence requires it. But that doesn’t mean having such an open mind that you consider everything everybody says a serious possibility.

### Implausible and Novel Claims can be True

A surprising or controversial or implausible claim may, of course, turn out to be true. When great breakthroughs in science, medicine, technology and other fields were first announced, they often seemed implausible. History presents a sorry record of discoveries the society of the time wasn’t ready to accept; for example, the Catholic Church required Galileo to renounce his claim that the Earth moved around the Sun.

It is extremely important that such ideas be given a fair and open hearing, and that they be accepted if the evidence supports them. But this does not mean that every time someone comes up with a new idea it should be taken as seriously as

ideas that are already supported by mountains of evidence. When the first vaccines were developed, the burden of proof was on those who developed them to show that they worked. In fact they shouldered this burden and provided evidence to back up their claims. But it would have been rash for someone to have gotten an injection just because someone in a lab coat handed them a needle syringe and offered the comforting words: “Well, you can’t prove that it won’t work.”

### Reserving Judgment

In many cases the evidence on both sides of an issue is inconclusive. In such cases, it is best to suspend belief, to refuse to conclude that either side is correct. For example, there is not much strong evidence on either side of the claim that there is life on other planets. Of course It may eventually find evidence that settles the matter (this would be much easier if there were extraterrestrial life and some of it showed up around here). But until then, it is reasonable to conclude that you just don’t have enough to go on and so you just don’t know.

We are often in no position to disprove a surprising or implausible claim. But it isn’t our responsibility to do so. Our failure to supply evidence *against* the claim does not somehow provide evidence *for* it. If someone else wants to convince others that a novel claim is true, it is up to them to provide evidence *for it*.

### Burden and Proof and the Law

An extremely important part of our legal system is that a defendant is presumed innocent until proven guilty. This means that the burden of proof is not on the defendant to show that he or she is innocent. The burden is on the prosecutor to show that the defendant is guilty. This makes sense, because the burden of proof is on the person who makes a claim (the claim here being that the defendant is guilty) rather than on the person on the other side.

To say that the defendant is presumed innocent is just another way of saying that we can’t use an appeal to ignorance to convict someone. We can’t argue: “well, she can’t show she didn’t do it, e.g., she doesn’t have an alibi, so she must be guilty.”

We typically have a higher standard for the burden of proof in the courtroom than in daily life, because the costs of mistakes are so high. In many cases, especially criminal cases, we require that the evidence be clear beyond any reasonable shadow of a doubt. This is so because the costs of mistakes are so high. Because the burden of proof lies with the state, a defendant is under no compulsion to testify at all. But juries, being human, often see failure to testify as some indication that the defendant is guilty.

### Exercises

1. Appeals to ignorance do not just arise with the Great Pumpkin. It is common for companies to argue that it has yet to be demonstrated that certain things are dangerous (e.g., smoking cigarettes, nuclear power plants, certain sorts of landfills and toxic waste dumps), and so we should continue to build or manufacture such things. Do such cases involve a fallacious appeal to ignorance (the answer may be different in different cases)?
2. Conspiracy theories often trade on appeals to ignorance. Since investigators haven't been able to show that something isn't the case, it is suggested, there is strong reason to think that it is. Give an example of this (it can be one you have read about or one you invented).
3. Appeals to authority are sometimes legitimate. Can you think of any special circumstances where an appeal to ignorance might be legitimate?
4. In the people on both sides of an issue are making claims that aren't at all obvious, then the burden of proof falls on both of them. Give an example of this sort.

And then, of course, we have these two claims:

5. It's reasonable to conclude that God exists. After all, no one has ever shown that there isn't a God.
6. It's reasonable to conclude that God does not exist. After all, no one has ever shown that there is a God.

## 10.6 Suppressed (or Neglected) Evidence

We commit the fallacy of **suppressed** or **neglected evidence** when we fail to consider (or simply overlook) evidence that is likely to be *relevant* to an argument. In this case we may include premises that are relevant, but we commit a fallacy because we leave out other information that is also relevant. Of course we cannot look at all of the evidence that might conceivably be relevant (that would be an endless task). But we should never neglect evidence that we know about or evidence that seems quite likely to bear on the issue.

The fallacy of suppressed (or neglected) evidence is a generic, catch-all label. So even where none of the names of fallacies we have learned seems quite appropriate, remember that we need to consider as much of the relevant evidence as we can when we evaluate an argument.

One of the easiest ways to make our position look good and to make alternative positions look weak is to suppress evidence that tells against our view as well as evidence that supports the alternative. So it is not surprising that relevant evidence is constantly suppressed in partisan disputes. In the courtroom lawyers only present evidence that will make their own case look good. In advertisements only one side of the picture is presented. In political debates, candidates almost always present only one side of the issue. In discussions about public policy, the partisans on each side of the issue often cite only those statistics that support their side of the case.

In many cases it would be expecting too much to think that those engaged in intense debates over such matters would present both sides of the issue in a fair and even-handed way. But there is often a big difference between winning an argument and thinking clearly, and if *we* must make up our own minds on the matter (as we must when we serve on a jury or vote in an election), we must consider as much of the evidence on each side of the issue as we can. And (as we noted when discussing the strawman fallacy), we cannot rely on the critics of a view to state it fairly, especially when the point at issue is highly controversial or arouses intense emotions.

Remember not to simply invoke the fallacy of suppressed (or neglected) evidence each time you see an argument that overlooks something that seems likely to be relevant. If the argument commits one of the fallacies we have studied, it is important to note that fact. And it is even more important to explain in some detail why the argument is weak. We will study this fallacy in greater detail in a later chapter on sampling.

## 10.7 Chapter Exercises

The task in this exercise set is to spot any of the fallacies we have studied thus far. To make things more interesting, it may be that some passages do not commit any fallacy at all. Identify any fallacies by name, then explain in your own words and in detail what is wrong with the reasoning in those cases where it is bad. Answers to selected exercises are found below.

1. Before the elections last November, some Democrats argued that the Republican revolution had to be stopped. After all, they said, the Republicans want to phase out health care for the elderly and on top of that, they want to take school lunches away from kids.

2 Before the elections last November, some Democrats claimed that the Republican revolution had to be stopped. After all, they said, almost all of the Republicans in Congress are rich and they only care about making things better for themselves.

3. If twenty is greater than ten, it is certainly greater than eight and twenty is greater than ten.
4. We really do have to conclude that O. J. Simpson is guilty. After all, he couldn't provide any alibi for where he was when the crime was committed—he claims he was home asleep, but no one can verify his claim—and the defense never suggested who else might have committed the crime if the defendant didn't.
5. Those who want to allow homosexuals to serve in the military forces hold an unreasonable and unacceptable position. They want gays to be automatically accepted into the military, and want no restrictions on sexual relations among homosexuals, whether on or off duty.
6. There is no way Bill Clinton could have been involved in the White Water scandal. If he had been, someone would have been able to prove it by now.
7. From a letter condemning hunting: “Please . . . don't try to ennoble your psychotic behavior by claiming that you are trimming the herd for the benefit of the herd.”
8. Champions of campaign finance reform have a hopeless view. They seem to think that if there were limits on contributions to political candidates there would be no more corruption in politics and the poor could also afford to run for office. But we will always have corruption, and it will always be easier for the rich to get elected.
9. Who are you to say that I drink too much? You aren't exactly Ms. Sobriety yourself.

### Answers to Selected Exercises

1. *Strawman fallacy* The Democrats in question were not personally attacking Republicans, so it is not an ad hominem fallacy. But their characterization of the Republicans' position is unfair; it distorts their view in a way that makes it appear much more frightening, and hence easier to attack.
- 2 *Ad hominem fallacy*. Here the Democrats in question were attacking the Republicans themselves, rather than their arguments.
- 3 This argument *affirms the antecedent*, and so it is valid. Both premises are true, so it is also sound. It is not a very interesting argument, but there is nothing substantive that is wrong with it. It's thrown in here just to remind you of fallacies we saw near the beginning of the course.

4. Taken *just by itself*, this argument commits the *fallacy of appeal to ignorance*. The mere fact that Simpson can't show that he is innocent does not show that he is guilty. In a larger context, where we know of positive evidence against Simpson, the lack of alibi does become a problem for him.
5. *Strawman fallacy*. This passage is an unfair, distorted characterization of the views of most people who think that homosexuals should be allowed to serve in the armed forces. No one in their right mind thinks that soldiers, whether they are homosexual or heterosexual, should have no restrictions on sex even when they are on duty (e.g., flying a combat mission).
6. Taken *just by itself*, this commits the *fallacy of appeal to ignorance*. The fact that evidence hasn't turned up doesn't mean that it won't. But in a larger context, *if* people have been looking very hard in the places where evidence would be (if it exists at all) and they have come up empty, it gives some support to the claim that he is innocent.
7. *Attack on the person (ad hominem)*
8. *Strawman*, combined with an *Either/or fallacy*. Why?



## Chapter 11

# Fallacies: Common Ways of Reasoning Badly

**Overview** Bad reasoning is said to be *fallacious*. Some fallacies are simple to spot, but others give the appearance of being good arguments, and it is easy to be taken in by them. In this chapter we learn about three more fallacies: begging the question, the either/or fallacy and, quite briefly, the fallacy of the line. We will also note problems involving inconsistency.

### Contents

---

<b>11.1 Begging the Question</b> . . . . .	<b>196</b>
<b>11.2 The Either/Or Fallacy</b> . . . . .	<b>203</b>
11.2.1 Clashes of Values . . . . .	206
11.2.2 Safeguards . . . . .	208
<b>11.3 Drawing the Line</b> . . . . .	<b>212</b>
<b>11.4 Inconsistency</b> . . . . .	<b>213</b>
<b>11.5 Chapter Exercises</b> . . . . .	<b>214</b>
<b>11.6 Summary of Fallacies</b> . . . . .	<b>217</b>

---

### Fallacies

Reasoning can go wrong in countless ways, but some of them turn up so often that it is useful to have name for them. This helps us spot them and to avoid them in

our own thinking. In this chapter we will study four of the most common ways reasoning and argumentation can go awry—four common fallacies—and discuss some ways to spot, and to avoid, them.

## 11.1 Begging the Question

You and Wilbur are chatting in the dorm. He's a rabid UCLA fan and is telling you how well their football team will do this season. But you are skeptical. In exasperation, he works the following argument into your conversation.

*You:* I just don't see UCLA having a winning season in football. They weren't very good last year, and they lost a some really good players.

*Wilbur:* You're wrong, and I can show it. UCLA will have a winning season.

*You:* Why?

*Wilbur:* Because they will.

When we put Wilbur's argument in standard form it looks like this:

Premise: UCLA will have a winning season.

Conclusion: Therefore UCLA will have a winning season.

Are you convinced? Of course not—but why not? What's wrong with the argument? Let's begin by asking whether the argument has some of the good features we've learned about so far.

**Relevance:** Is the premise relevant to the conclusion? Well, if the premise is true, so is the conclusion, so the premise could scarcely be more relevant.

**Validity:** Is the argument valid? Well, *if* the premise is true, then the conclusion has to be true (since the premise and the conclusion are the same). So the argument is valid.

**Soundness:** Is the argument sound? That depends on whether the premise is true or not. We may not be sure about that, but we can see that there is something wrong with this argument even if the premise is true by devising a similar argument with a true premise..

- Suppose you aren't sure who the U.S. President was in 1922 and Wilbur argues: The President in 1922 was Warren G. Harding; therefore the President in 1922 was Warren G. Harding. Here the premise is true,

so the argument is sound. But if you were in doubt before, should this argument convince you that Harding was President in 1922? This argument is no better than Wilbur's first argument.

The premises of the winning-season argument are relevant to the conclusion, the argument is valid, and it might even be sound. Still, something is very wrong with it in the context of your discussion with Wilbur. What is it (try to answer this question before going on)?

### Using Arguments to Convince

Arguments have various purposes. Sometimes we just use an argument to see what follows from what (if I spend \$50 on CDs this week, how much will that leave me for food?). But one of the main functions of argument is to convince someone of something that they do not already believe (if they already accept our claim, there is no point trying to convince them).

When we try to legitimately convince someone with an argument, we are trying to get them to accept a claim (the conclusion) by giving them reasons (premises) to believe it. But if we are to accomplish this, we must use premises that the other person accepts. After all, if we use premises that they don't accept, then even if our argument is deductively valid, there is no reason for them to believe that our conclusion is true.

What counts as a legitimate premise depends on the context. If you are arguing with someone who already believes that democrats always make better public officials than republicans, then you could use the claim that we should have a democratic majority in the United States Senate in arguments with that person.

You both agree on this premise, it is common ground, so it is quite reasonable to use it in this context. But if you are arguing with someone who doesn't believe that democrats always make better officials, you cannot use this premise. If the other person doesn't think your premise is true, then no matter how elaborate your argument, she won't have any reason to accept your conclusion.

So it is fair to assume different things in different contexts. But if someone doubts that your view is true, it will never be appropriate to use your view as a premise to convince them that your conclusion is true. No matter what the context, if you doubt that OU will have a winning football season, Wilbur cannot reasonably use the claim that they will have a winning season as a premise in arguing for his conclusion that they will have a winning season.

Of course no one would be taken in by Wilbur's argument, but we will see that it is often possible to give the sort of argument that Wilbur gave, but to *disguise*, it so that its defects are much harder to spot.

**Begging the question:**  
assuming what you are  
trying to prove

**Begging the Question** We commit the *fallacy of begging the question* when we assume the very thing as a premise that we're trying to show in our conclusion. We just assume the very thing that is up for grabs. This is a fallacy, because if a certain point is in dispute, we cannot fairly assume it in our discussion.

Let's continue for just a moment with our a blatant example where the problem is obvious. We will then and work our way up to more difficult cases.

Premise: UCLA will have a winning season.

Conclusion: Therefore UCLA will have a winning season.

If you were in doubt about the conclusion, you won't accept the premise, and so you won't be persuaded by this argument. After all, the reason Wilbur gives you for accepting the conclusion is the very point at issue in your discussion.

Of course no one is fooled by such an obviously bad argument. But now suppose that Burt restates his conclusion. If he rephrases it in different words, then the flaw may be harder to detect. But before looking at examples, we should note a very important point that emerges from our discussion thus far.

Arguments that beg the question have premises that are relevant to their conclusions, they are deductively valid, and many of them are even sound. For example, the argument "Two is an even number; therefore, two is an even number," is sound. Does this mean that relevance, validity, and soundness doesn't matter after all? No. It merely shows that in some contexts there are additional things that matter in an argument.

### **Begging the Question by Rephrasing the Conclusion**

Sometimes people rephrase the conclusion (put it in different words), and then use the result as a premise. This can be confusing if they use technical jargon or puts things in a long-winded way. Consider the following dialogue:

*Albert:* One of these days I think we'll have a successful communist country.

*Burt:* Communism will never succeed, because a system in which everything is owned in common will never work.

In standard form Burt's argument looks like this:

Premise: A system in which everything is owned in common won't work.

Conclusion: Therefore, communism will never work.

The conclusion of this argument may well be true. But the very point at issue between Albert and Burt is whether or not communism could work. And in this context, Burt's premise cannot be used to support his conclusion, since it simply restates the conclusion in different words. To see this, note that

- communism = a system in which everything is owned in common, and
- will never succeed = will never work

This argument is really no better than the argument that OU will win because they will win. But when the premise restates the conclusion in different words, the fallacy can be harder to detect.

Here is another example.

*Al:* I don't know about all the things in the *Bible*. Like the song says, it ain't necessarily so.

*Burt:* That's wrong. The *Bible* is the word of God.

*Al:* How in the world do you know that?

*Burt:* Well, the *Bible* says that it's the words of God, and it's divinely inspired.

In standard form:

Premise 1: The *Bible* says that it's the word of God.

Premise 2: The *Bible* is divinely inspired.

Conclusion: Hence, the *Bible* is the word of God.

Al and Burt will agree that the first premise of the argument is true. But since the point at issue here is whether the *Bible* is the word of God, Burt's second premise begs the question. In order to view Premise 2 as evidence for the existence of God, you already have to *believe* in the existence of God.

Put another way, anyone (like Al) who doubts that the *Bible* is the word of God will also doubt whether it is divinely inspired. Since being divinely inspired and being inspired by God are the same thing, the two claims here say virtually the same thing. This doesn't mean that Burt's conclusion is false. But if the goal is to convince Al, then Burt needs *independent* support for his claim. He could either

1. Employ some other premise Al will accept, or
2. Defend premise 2 (using premises Al will accept).

Here is another argument that suffers from the same malady: "Democracy is the best form of government, since the best system is one in which we have government by the people." Put this argument into standard form and analyze it.

### **Begging the Question by Generalizing the Conclusion**

A trickier form of begging the question arises if we *generalize* the conclusion and use the result as a premise. This problem is illustrated in the following dialogue:

*Beth:* There's nothing wrong with a couple of cold beers on a hot summer day.

*Seth:* Oh no. Drinking beer is wrong!

*Beth:* Why in the world is that?

*Seth:* Well, because drinking alcohol is wrong.

In standard form Seth's argument looks like this:

Premise: Drinking alcohol is wrong.

Conclusion: Therefore, drinking beer is wrong.

Here Seth's premise does not simply restate his conclusion. But it does generalize it (since beer is one type of alcohol). Beth surely knows that beer is one variety of alcohol, and since the point at issue is whether drinking beer is wrong, she won't accept the more general claim that drinking alcohol is wrong.

### **Begging the Question in more Subtle Ways**

It is also possible to beg the question in more subtle ways. For example:

*Alice:* I know abortion is a terrible thing, but I don't think it should be illegal.

*Beth:* But you're overlooking the fundamental point. Abortion is murder. And we certainly should have laws against that. So we should have laws against abortion.

In standard form Beth's argument looks like this:

Premise 1: Abortion is murder.

Premise 2: We should have laws against murder.

Conclusion: So we should have laws against abortion.

We can assume that Alice and Beth and virtually everyone else agrees that we should have laws against murder, so it is perfectly appropriate for Beth to assume this as her second premise. But the point at issue is really just whether abortion is wrong in a way that would justify having a law against it. Alice denies that it is, so she certainly wouldn't accept Beth's premise that abortion is murder. In this context, Beth's first premise assumes the point at issue, and so begs the question.

Of course Beth's first premise might be true. But since Alice would not accept it at this stage of their discussion, Beth needs to give some further argument to support it. If she can do this, then Alice will almost certainly accept Beth's further claim that we need laws against abortion.

Some people draw a distinction between the fallacy of *begging the question* and the fallacy of *circular reasoning*. We needn't worry about such fine distinctions here, though, so we'll use these two labels interchangeably to stand for the same fallacy.

### Question Begging Labels

Sometimes we characterize a view or group or person in a loaded way, with a label that begs the question against it (or them or her). This happens when the label only makes sense if the view or group or person is defective in some way. For example, labels like 'redneck' and 'welfare queen' suggest that members of a certain group are guilty of certain practices or have certain dangerous views.

It is also possible to use labels to beg the question in favor of a position or group or person. For example, labels like 'The moral majority' suggests that the group's views represent those of the majority and that they are right. Perhaps they do, but you can't use a label to settle the question.

Of course no one would think that a label completely settles the matter, but labels do predispose us to think about things in certain ways.

### Exercises

1. In each of the following determine whether a question is begged. If it is, say as precisely as you can just how the fallacy is committed (Does the premise restate the conclusion? Does it just generalize it? Is something more subtle going on?).

1. All freshmen at OU should have to attend computer orientation, because all university students should go to such an orientation.
2. All freshmen at OU should have to attend computer orientation, because without it they won't be prepared to use computers, and they'll have to master that skill to have a decent job in the coming century.
3. Capital punishment is morally wrong, because the *Bible* says that it is.
4. The belief in God is nearly universal, because nearly everybody, in every culture and every historical period, has believed in God.
5. It is important that we require handguns to be registered, because it keeps guns out of the hands of children (they are too young to be registered) and dangerous criminals (since a background check is required for registration).
6. It's important that we require handguns to be registered, since we need some sort of record who the owners of fire arms.

7. Capital punishment is morally wrong, because it's always wrong to take the life of another person.
  8. God exists, because nearly everybody, in every culture and every historical period, has believed in God.
2. Give an example of a question-begging argument where one of the premises simply restates the conclusion.
  3. Give an example of a question-begging argument where one of the premises simply generalizes the conclusion.
  4. What sorts of premises is it legitimate to assume in a given context? How can you determine whether a premise is legitimate or not?

### Answers to Selected Exercises

1. In each of the following determine whether a question is begged. If it is, say as precisely as you can just how the fallacy is committed (Does the premise restate the conclusion? Does it just generalize it? Is something more subtle going on?).

- 1.1 All freshmen at OU should have to attend computer orientation, because all university students should go to such an orientation.

This begs the question. The conclusion generalizes the premise, and anyone who doubted that OU freshmen should attend a computer orientation would (probably for the same reasons) doubt that all university students should.

- 1.2 All freshmen at OU should have to attend computer orientation, because without it they won't be prepared to use computers, and they'll have to master that skill to have a decent job in the coming century.

The premise offers an independent reason why freshmen here should go through computer orientation. Perhaps the argument has flaws, but it doesn't beg the question.

- 1.3 Capital punishment is morally wrong, because the *Bible* says that it is.

This argument does not beg the question. Perhaps it has other flaws, but the premise does not restate or generalize or presuppose the conclusion. It tries to provide *independent* support for it.

- 1.4 God exists, because nearly everybody, in every culture and every historical period, has believed in God.



This argument does not beg the question. It offers a reason that is quite independent of the conclusion as support for the conclusion. The argument is actually a version of an argument from authority (an appeal to tradition and to what most people think).

## 11.2 The Either/Or Fallacy

### Disjunctions

Earlier we encountered a type of compound statement called a *conditional*. In this module we will be concerned with a second kind of compound sentence; it's called a *disjunction*. **Disjunction:** an either/or sentence

A **disjunction** is an “either/or” sentence. It claims that (at least) one *or* the other of two alternatives is the case. For example:

1. Either the butler did it, or the witness for the defense is lying.
2. Either I have a throat infection, or I have the flu.
3. “ ‘Hey,’ I said, ‘When you [write novels], do you sort of make it up, or is it just, you know, like what happens?’” [Martin Amis, *Money*].
4. We will either balance the federal budget this year, or we will stand by and watch our country go broke sometime in the next quarter century.
5. Either the dead simply cease to exist and have no perceptions of anything, or else they go on to a better life after death [Socrates, from Plato’s dialogue, the *Apology*].

The two simpler sentences that make up a disjunction are called **disjuncts**. The order of the disjuncts in a disjunction doesn’t matter (you can reverse their order, and the resulting disjunction will mean the same thing as the original disjunction).

The first sentence in the list above says that either the butler did it *or* the witness for the defense is lying. So it will be true if the butler did do it *or* if the witness is lying. And it will be false if both of these disjuncts is false.

### The Either/Or Fallacy

We commit the **either/or fallacy** when we assume that there are only two alternatives when in fact there are more. We commit this fallacy when we mistakenly suppose that a disjunction is true when it actually is false. The either/or fallacy gets its name from the fact that we act as though *either* the one alternative is true **Either/Or fallacy:** assuming there are fewer alternatives than there actually are

or else the other alternative is, though in fact there are more than just these two alternatives. In such cases we have overlooked some third alternative. For example in sentence 1. (about the butler), we may have overlooked the possibility that the witness made a honest mistake (maybe her eyesight isn't what it used to be).

The either/or fallacy goes by a variety of names. It is sometimes called the *false-dilemma fallacy*, the *black and white fallacy*, the *fallacy of false alternatives*. It often results from what is called *all-or-none thinking*.

These names reflect the nature of the fallacy. For example, talk of black or white thinking suggests a tendency to think in *extremes*, to see things as definitely one way *or else* definitely the other, without any room in between for various shades of gray. And talk of all-or-none thinking suggests a tendency to think that things must be all one way or all another (when in fact the truth may like somewhere in the middle). Let's see how some of the sentences in the list above might involve the either/or fallacy.

1. Either the butler did it, or the witness for the defense is lying.

As we noted above, the butler may be innocent and the witness may simply be mistaken for some reason or another.

2. Either I have a throat infection or I have the flu.

Perhaps I have correctly narrowed the possibilities down to these two, in which case the disjunction is true. But I may have jumped to the conclusion that I have one or the other of these infirmities, when in fact it is only my allergies acting up again (perhaps the pollen count is high lately).

3. “ ‘Hey,’ I said, ‘When you [write novels], do you sort of make it up, or is it just, you know, like what happens?’” [Novelist] “Neither” [Martin Amis, *Money*, p. 87].

Here John Self, a character in Martin Amis's novel *Money* asks a novelist whether he just makes everything up or whether it he writes about things that have really happened. The novelist replies that he doesn't do either of the two. Real-life events give him lots of ideas, but he is constantly changing them in his imagination as he writes. The novelist is pointing out, in his one-word response (“Neither”), that Self is committing the either/or fallacy.

4. Either the dead simply cease to exist and have no perceptions of anything, or else they have a good life after death.

Socrates asserts this disjunction in the course of an argument to the conclusion that death is nothing to fear. He seems to overlook the possibility that there is life after death, but that it will be unpleasant.

### Examples of the Either/Or Fallacy

Since the either/or fallacy makes things look simpler than they really are, it makes for pithy, memorable slogans. You will be familiar with some of the following examples:

- America: Love it or Leave it.
- You are either part of the solution or you are part of the problem.
- You are either with us or against us.

Another common pair of slogans several decades ago were

- “Better dead than red” (a favorite of proponents of the arms race at the height of the cold war), and
- “Better red than dead” (a favorite of their opponents).

Both of these slogans rested on the assumption that there were really only two alternatives: either we engage in an escalating arms race with the Soviet Union or they will crush us. Such claims often do make a point, but snappy slogans by themselves rarely make for good reasoning.

### More Complex Examples of the Either/Or Fallacy

**More Complicated Sentences** Many examples of the either/or fallacy can be more difficult to spot. For example, one frequently hears the following sorts of claims in current debates over policy issues.

1. Either we keep teaching the Western Canon (the great literary and philosophical works of the Western world), or we just let each professor teach whatever junk he wants to.
2. We either have to institute the death penalty, or we will have to live with the same people committing terrible crimes over and over (each time they are released from prison).

With a little thought we can see that there are more than two alternatives in each case. But when we hear such claims in conversation, they often go by so fast,

and may be asserted with such confidence, that we don't realize how much they oversimplify the situation.

The either/or fallacy is often committed along the strawman fallacy. For when we simplify someone's views into two, easily-attacked alternatives, we typically substitute a weakened version of their view for the view that they really hold.

### 11.2.1 Clashes of Values

Many of the difficult moral and political issues of our day involve clashes of values. Virtually all of us think that all of the following are important. Consider the following values:

1. Freedom or liberty
2. Security
3. Majoritarianism (majority rule, i.e., democracy)
4. Equal opportunity
5. Community moral standards

For example, virtually all of us believe that democracy, the rule of the many, is a good thing. But democracy can be in tension with other values, most obviously individual rights and liberties. A majority can tyrannize a minority just as much as a dictator can. For example, until the 1960s poll taxes and other public policies made it almost impossible for African-Americans in many parts of this country to vote.

Freedoms or liberties or rights are in tension with other values. For example, allowing people a great deal of freedom can make our lives less secure in various ways. Indeed, the various types of liberty and freedom can even be in tension among themselves. Unlimited freedom of the press may infringe on a person's right to have a fair trial, for example, or violate people's rights to privacy.

Most of us think that it is also important to make sure that everyone in our country (especially children) have a decent standard of living (enough food to survive, basic medical care, etc.). But the two values, freedom to spend one's money as they like and helping others, are in some tension, because in our society the only way to ensure that everyone has a reasonable standard of living is to tax people and redistribute the money to people who don't have much (in the form of food stamps, welfare payments, etc.).

In such cases, it may be tempting to pose the issues in terms of two stark alternatives: freedom (to keep what one owns) vs. everyone having a reasonable standard of living. These cases are difficult, because both values are important to most of us. But there is a definite tension between them, and finding satisfactory ways to reconcile them is difficult.

### Disjunctions in the Guise of Conditionals

Disjunctions can be restated as conditionals. For example, we can restate claim “Either you are part of the solution or you are part of the problem” as the conditional

- If you are not part of the solution, then you are part of the problem.

And we can restate the claim “Either I have a throat infection or I have the flu,” as the conditional

- If I don’t have a throat infection, then I have the flu.

If I assert either of these conditionals, I may be just as guilty of the either/or fallacy as if I had asserted the original disjunction, but it is harder to spot this fallacy when we are dealing with a conditional. If you come across the claim

- If we don’t balance the federal budget this year, then we will stand by and watch our country go broke sometime in the next twenty five years

It may be far from obvious that this passage involves the either/or fallacy. To see that it does, we must see that it is just another way of saying “We will either balance the federal budget this year, or we will stand by and watch our nation go broke sometime in the next twenty five years.”

Finally, we should note that the either/or fallacy could be involved if someone said that there were only three alternatives when in fact there were four. Indeed, the fallacy is committed whenever someone claims that there are fewer alternatives than there actually are.

### Why it’s Easy to Commit the Either/Or Fallacy

1. It takes less energy and imagination to suppose that there are only two alternatives than to try to figure out whether there are additional possibilities.
2. Language is full of simple opposites—good vs. bad, right vs. wrong, us vs. them—and this can encourage us to think in oversimplified terms.
3. It is easier to persuade others to accept our view about something if we can convince them that the only alternative is some very extreme position. So characterizing the situation in terms of a limited number of options often makes it easier to defend our position.

4. Prejudices and stereotypes can make it easier to think in all-or-none fashion (we will return to this in a later chapter). For example, extremists of various sorts tend to see issues in very simple terms; it is either us or them. For this reason, such people are often unwilling to compromise.
5. There are two kinds of deductively valid arguments that involve disjunctions, and below we will see how these sometimes encourage us to commit the either/or fallacy.

### 11.2.2 Safeguards

In any particular case there may really only be two alternatives, and there are certainly issues on which we should not be willing to compromise. But we shouldn't assume this without considering the matter. Here are some safeguards against the either/or fallacy.

1. When you encounter a disjunction (or a conditional) consider the possibilities. Has the other person overlooked some genuine alternatives?
2. Avoid the temptation to think in extremes. Difficult issues rarely have simple solutions, so we need to at least consider a range of options.
3. Be especially wary if someone argues that the only alternative to their position is some crazy-sounding, extreme view.

### Arguments involving Disjunctions

There are two kinds of arguments that involve disjunctions. They are relevant here, because when we use them, it can be easy to commit the either/or fallacy.

In Chapter 3 we studied conditional arguments. Here we will learn about two kinds of *disjunctive arguments*, two important argument forms or formats that involve disjunctions.

### Disjunctive Syllogism

#### Disjunctive Syllogism:

*A* or *B*

Not-*A*

So *B*

1. Either the butler did it or the witness is lying.
  2. The witness isn't lying (she's as honest as the day is long).
- So (3) The butler did it.

Arguments having this form are called **disjunctive syllogisms**. *All* arguments with this form are deductively valid. They involve a simple process of elimination; one premise says that there are only two possibilities and the second premise eliminates one of the two. This leaves only one possibility as the conclusion.

Disjunctive syllogisms have the form:

1. Either *A* or *B*.
2. But *A* is not true.
- So (3) *B* is true.

### Constructive Dilemmas

Socrates gives the following argument for the conclusion that death is nothing to fear (Plato reports the argument in his dialogue the *Apology*).

1. Either the dead simply cease to exist and have no perceptions of anything, or else they go on to a better life after death
2. If the dead simply cease to exist, then death is nothing to fear [it would be like a long, restful sleep].
3. If the dead go on to a better life, then death is nothing to fear.
- So (4) Death is nothing to fear.

Arguments having this form are called **Constructive Dilemmas**. *All* arguments with this form are deductively valid.

In the first premise we narrow the range of alternatives down to two. Then, even if we don't know which of the two is the case, we claim (in premise 2) that if the first alternative is true, then such and such follows. We repeat this strategy, adding (in premise 3) that if the second alternative is true, then (the same) such and such follows. So if *either* of them is, such and such must follow (in this case, that death is nothing to fear).

Constructive dilemmas have the form:

1. Either *A* or *B*.
2. If *A*, then *C*.
3. If *B*, then *C*.
- So (4) *C*.

#### Constructive Dilemma:

*A* or *B*  
 If *A* then *C*  
If *B* then *C*  
 So *C*

### Disjunctive Arguments and the Either/Or Fallacy

When we encounter either type of disjunctive argument, we should ask whether its disjunctive premise is true or whether we have an either/or fallacy. Such fallacies

are especially easy to overlook in such contexts because the argument may be good in several ways that lead us to overlook the false disjunctions. In particular,

1. Both sort of arguments are always deductively valid so the formal reasoning is correct.
2. All of the premises other than the disjunction may well be true.
3. The person giving the argument may spend a lot of time defending the non-disjunctive premises. This may focus our attention on them, leading us to overlook potential problems with the disjunction. In Socrates' argument about death, for example, a lot of time might be spent defending the claim that an eternal sleep is not to be feared, and this may lead us to overlook problems with the first, disjunctive, premise.

### Exercises

When you encounter a disjunction (or a conditional) it is always worth asking whether it commits the either/or fallacy. Thought is required; we want to be alert to the possibility that this fallacy has been committed, but we don't want to jump too quickly to the conclusion that it has.

In each of the following passages determine whether the either/or fallacy has been committed or not. In those cases where it has (a) say precisely how it has been committed, and (b) explain what might be done to strengthen the argument so that it doesn't commit this fallacy.

1. Mother to son: "Are you going to college, or are you going to be a bum like the Jones boy?"
2. "Hallmark, when you care enough to send the very best."
3. Either a positive integer is even or else it is odd.
4. Roseanne: "How bad is it? I mean, are we going to have to eat cat food, or just the kids?"
5. If you can't beat'em, join'em.
6. If God doesn't exist, then anything is permitted.
7. We obviously cannot legalize drugs, as some people recommend. For if drugs aren't illegal, we will be encouraging people to use them. And it is clear that drugs are extremely dangerous. So it's better to live with the current situation than to try to change things in this extreme sort of way.



8. We will either balance the federal budget this year, or we will stand by and watch our country go broke sometime in the next quarter century.
9. Either there is a God, or there isn't.
10. Either the dead simply cease to exist and have no perceptions of anything, or else they go on to a better place after death.
11. Either we keep teaching the Western Canon (the great literary and philosophical works or the Western world), or we just let each professor teach whatever garbage he wants to.
12. Either we have to institute the death penalty, or we will have to live with the same people committing terrible crimes over and over (each time they are released from prison).
13. Polling questions and opinion surveys often require you to select from a fairly restricted set of alternatives (e.g., should we increase defense spending or should we lower it?). Give examples, either ones that you have encountered or ones that you construct, which illustrate this.

### Answers to Selected Exercises

In each of the following passages determine whether the either/or fallacy has been committed or not. In those cases where it has (a) say precisely how it has been committed, and (b) explain what might be done to strengthen the argument so that it doesn't commit this fallacy.

1. Mother to son: "Are you going to college, or are you going to be a bum like the Jones boy?"

OK., so the first one is too easy.

2. "Hallmark, when you care enough to send the very best."

This is a nice advertising hook. The conditional, "When you care enough to send the very best, you send Hallmark," is equivalent to the disjunction "Either you don't care enough to send the best, or you send Hallmark." So if you send any card that isn't a Hallmark you really just don't care much about the person you send it to (you louse). Commits the either/or fallacy.

3. Either a positive integer is even or else it is odd.

This does not commit the either/or fallacy. There really are only two alternatives here. A positive integer must be one or the other. The claim here is perfectly true, and it doesn't involve any fallacy whatsoever.

4. Roseanne: "How bad is it? I mean, are we going to have to eat catfood, or just the kids?"

5. If you can't beat'em, join'em.

The key to working this is to note that the conditional here is equivalent to the disjunction:

Either you beat them or you join them.

### 11.3 Drawing the Line

When there are borderline cases, cases where we can't be sure whether a word applies to something or not, the word is **vague**. Vague words are not completely precise. For example, the word 'bald' is vague. There are many people who have a bit of hair, and they are not clearly bald or clearly non-bald. Again, some things are clearly red and some are clearly not red, but at the edges (near shades of orange and shades of purple), there are unclear cases.

When we encounter a vague word, there is usually no need to try to make it precise (and in fact any way of doing so would be somewhat arbitrary). Indeed, although precision is often desirable, it is actually a good thing that many of our words are vague. If we had to learn exactly how many hairs someone had to have to be non-bald, or precisely what shades counted as red, we could never learn to use these words. In fact, most of the adjectives (and some of the other words) in our language are vague.

It is not uncommon to hear people argue as follows:

You can not draw a definite line between  $X$  and  $Y$ , so there really is no difference between  $X$ s and  $Y$ s.

We will call this the **line-drawing fallacy**. There are a great many cases where we can not draw a definite, non-arbitrary, line between two things, but there are still many clear cases of  $X$ s and  $Y$ s (even though there are borderline cases as well).

For example, it is true that we cannot draw a line that neatly separates each person into either the group of people who are definitely bald or else into the group of people who are definitely not bald. There are borderline cases here. But this

doesn't mean that there are not clear cases of bald people (e.g. Telley Savalas) and clear cases of people who are not bald (e.g., Howard Stern). Again, it isn't possible to draw a precise line between day and night, but it is definitely day at 2:00P.M. and definitely night at 2:00A.M.

No one would agonize much over these examples, but there are other cases involving line-drawing that matter more. For example, you probably can not draw a precise, non-arbitrary line that will separate all weapons into those that should be legal (e.g., hunting rifles) and those that should not be (e.g., nuclear warheads), but this doesn't mean that there aren't clear cases of each.

Again, someone might argue that there is no way to draw a line between a fetus that is one day old and a fetus that is nine months old. Indeed, this could be part of an argument against abortion; we can imagine someone urging: "Abortion shouldn't be allowed, because there is no place where you can draw a line between the fetus's being a human being and its not being one." But there are certainly important differences between a very undeveloped fetus and a new-born infant.

Often we can't draw a precise, non-arbitrary line that separates everything neatly into either *Xs* or *Ys*. But that does not mean that there are not many completely clear cases of each. Moreover, it does not follow that just any place we draw the line is as good as any other. Any attempt to distinguish day from night that counts 1:00P.M. as night is simply wrong.

## 11.4 Inconsistency

The key thing to remember is that when some person (or group) is inconsistent, *at least one of the things they say must be false*.

People are rarely blatantly inconsistent. We don't often say something in one breath and then say the exact opposite in the next. But people do sometimes say one thing in one setting and then deny it in another. Politicians, for example, often tell different audiences what they believe the audiences want to hear, but all of us are susceptible to the temptation to do this.

People also sometimes promise to do several things that cannot all be done together. For example, someone running for office may promise to cut taxes, keep social security and medicare spending at their current levels, and beef up defense. It is unlikely that all three things can be done at once.

Organizations can also send inconsistent messages. For example, one member of a large organization may make an announcement in a way that allows other members to maintain "deniability" (i.e., that allows them to avoid taking responsibility for it).

To test your understanding, say what is wrong with the following argument

“Well, I agree that your argument is deductively valid and that all of its premises are true. But I still think that its conclusion is false, and who are you to say I’m wrong?”

We will return to issues involving inconsistency in Chapter 19, on cognitive dissonance.

## 11.5 Chapter Exercises

The task in this exercise set is to spot any of the fallacies we have studied thus far. To make things more interesting, it may be that some passages do not commit any fallacy at all. Identify any fallacies by name, then explain in your own words and in detail what is wrong with the reasoning in those cases where it is bad. Some fallacies from earlier chapters may appear here. Answers to selected exercises are found on page 215.

1. We can never bring complete peace to countries like Albania. There has been ethnic strife there for centuries. We can’t undo all that damage. So we should just stay out.
2. *Alice*: Burt and I both endorse this idea of allowing prayer in public schools, don’t we Burt?  
*Burt*: I never said any such thing.  
*Alice*: Hey, I didn’t know you were one of those atheist types.
3. This is one argument I’m going to win. My point is very simple. Either OU will win their division of the Big 12 outright (outright means that they win it without a tie) or else they won’t.”
4. If the prosecution fails to prove beyond a reasonable doubt that the defendant is guilty, then we ought to find the defendant not guilty. So we should indeed return a verdict of not guilty, since the prosecution has failed to offer convincing proof of the defendant’s guilt.
5. Either we allow abortion or we force children to be raised by parents who don’t want them.
6. We must not legalize marijuana. The legalization of marijuana would mean that it would not be a criminal act to possess marijuana. But certainly it is, and must remain, a criminal act to possess an illegal drug like marijuana. So we should oppose legalizing it.

7. The views of those who favor the mandatory use of seat belts are ridiculous. They claim that if everyone is required by law to wear seat belts when riding in a car, then there will be no more automobile fatalities, and that serious automobile injuries will be almost entirely eliminated. But that is an ridiculous view. For clearly some auto accidents are so bad than even the best seat belts would not prevent injury or death. So it is silly to have a law requiring the use of seat belts.

8. Almost every advertisement you see is obviously designed, in some way or another, to fool the customer: the print that they don't want you to read is small; the statements are written in an obscure way. It is obvious to anybody that the product is not being presented in a scientific and balanced way. Therefore, in the selling business, there's a lack of integrity. — Richard P. Feynman

9. *Burt*: Unless we construct a dam and a power plant in this area within the next ten years, we won't be able to meet the significantly growing demand for electrical power.

*Wilbur*: What you're saying is that you couldn't care less what happens to the plant life and wildlife in this area or even to human lives that might be dislocated by building the dam.

10. When you are buying a new car battery, it's hard to know which car battery you should choose. But remember one thing: Chuck Yeager says that AC Delco batteries are the best you can buy. And Chuck Yeager is one of the greatest test pilots of all time.

11. Abortion shouldn't be allowed, because there is no place where you can draw a line between the fetus's being a human being and its not being one.

12. Reporter to participant in a cow-chip throwing contest: "Why would anyone want to throw pieces of dried cow dung?"

Contestant: "Well, it beats the hell out of standing around holding them."

13. This piece of legislation is designed to exploit the poor. After all, it was written and sponsored by one of the richest people in the state.

### Answers to Selected Exercises

1. *Either/or fallacy* This passage presents us with only two alternatives: either we can bring complete peace or we should just stay out. It is fallacious because

it overlooks various intermediate possibilities. For example, we might be able to stop a good deal of the murder of innocent people, even if we can't stop all of it.

2. This passage doesn't contain a complete argument, but what Alice says does suggest that she is committing the *either/or fallacy*. She claims, in effect, that *either* you support prayer in public schools *or* else you are an atheist. Also probably attacks a strawman.

3. The speaker here asserts that a disjunction (and "either/or") sentence is true. In this case, there really are only two options: either OU wins their division outright or else they won't (if they tie for first, then they don't win outright). There is no either/or fallacy here.

4. *No fallacy*.

5. *Either/or fallacy*. Presents only two alternatives, when in fact there are several more.

## 11.6 Summary of Fallacies

**Fallacy of irrelevant reason (or irrelevant premise).** We commit this fallacy if we offer a premise to support a conclusion when the premise is irrelevant to the conclusion.

**Fallacy of Argument against the Person** An irrelevant attack a person, rather than on her position or argument. The fallacies Latin name, *ad hominem* (“against the person”) is still in common use.

**Strawman Fallacy** We commit this fallacy if we distort or weaken someone’s position or argument in an effort to discredit it. It is often tempting to do this, because it is much easier to attack a distorted version of a view than to attack the real thing. The chief safeguards here are to (1) be aware of the natural human tendency to characterize opposing views in a way that makes them easier to attack or dismiss, (2) discuss the strongest version of a view you don’t like, and (3) do not rely on the critics of a view to state it fairly.

**Suppressed (or Neglected) Evidence** We commit this fallacy if we fail to consider (or simply overlook) evidence that is likely to be relevant to an argument. Like the generic fallacy of irrelevant reasons, the fallacy of suppressed (or neglected) evidence is a generic, catch-all label.

**Begging the Question** Assuming (without argument) the very point that is up for grabs in a given discussion.

**Appeal to Ignorance** We commit this fallacy if we defend a view by pointing out that others can’t show that it’s false. The fact that they are ignorant (don’t know) of evidence that would show we are wrong does not mean we are right.

**Either/or Fallacy** We commit this fallacy if we assume that there are only two alternatives when in fact there are more. The chief safeguards are to (1) consider all the genuine alternatives in a given case, (2) avoid the temptation to think in extremes, and (3) be wary if someone urges that the only alternative to their view is some crazy-sounding extreme view.

**Fallacy of the Line** We commit this fallacy when we argue that because we cannot draw a definite, non-arbitrary line between two things, there really isn’t any difference between them.

**Inconsistency** The basic problem with an inconsistent set of claims is that at least one of them must be false. In this module we learned about several ways of camouflaging inconsistencies.





## **Part V**

# **Induction and Probability**



## Part V. Induction and Probability

Life is uncertain, and we often must act in cases where we can't be sure about the effects our actions will produce. Still, some things seem much more likely than others. In this module we will examine a range of cases where it is possible to measure how probable something is; we call the numbers used in such measurements *probabilities*.

In Chapter 12 we will get a quick overview of the general issues.

In Chapter 13 we will learn rules for calculating the probabilities of certain important types of sentences.

Finally, in Chapter 14 we will learn how to deal with conditionals probabilities and some related concepts.



# Chapter 12

## Induction in the Real World

The race is not always to the swift, nor the victory to the strong. But that's the way to bet.

—Damon Runyon

**Overview:** In this chapter we introduce the notion of probability and see how it effects every aspect of our lives.

### Contents

---

<b>12.1 Life is Uncertain</b> . . . . .	<b>223</b>
<b>12.2 Inductively Strong Arguments</b> . . . . .	<b>226</b>
<b>12.3 Chapter Exercises</b> . . . . .	<b>227</b>

---

### 12.1 Life is Uncertain

Almost everything in life is uncertain, so we can't help but deal with probabilities. We don't have any choice. Still, some things are much more likely than others. It is reasonably probable that there will be at least a little snow in Norman some time this year, less probable that there will be over six inches of snow this January, and very improbable that there will be any snow at all in June. In this module we will study the notion of probability.

Many of the examples in this module involve games of chance. We employ such examples because they are relatively simple, and many of you are familiar

with them already. But probability is important in many other settings, and once we learn how it works in these simpler cases, we will see that it turns up almost everywhere. Here are a few of the many cases where we must make important decisions on the basis of uncertain information.

**Visiting the Doctor:** Often a person's symptoms and test results are compatible with several different diagnoses, but some of the diagnoses may be more probable than others. Furthermore, the outcomes of many medical treatments are uncertain. You or a loved-one may one day have to consider the probability that a risky treatment will improve a very bad medical condition.

- Example: suppose that you have to weigh the odds of a type of operation in which 65% survive with a much higher quality of life—but 5% of the patients die in surgery.

**Selecting a Major:** A college degree will make it easier for you to get a job, but it requires a lot of time and money, and it won't guarantee a good job once you graduate. And once you do decide to attend college you have to select a major. Perhaps the fields you like best offer fewer job opportunities than fields you like less. How should you decide what to major in?

**Divorce:** There is a reasonable chance that if you get married you will end up getting divorced. According to the *Oklahoma Gazette* (Nov 20, 1997) Oklahoma has the second highest divorce rate in the nation (trailing only Nevada). A 1987 University of Wisconsin study based on 1987-89 data found that in the country overall, 27% of married couples divorced in the first decade after their marriage. The rate has declined since then, but this is only a ten year period. Of course few people think that they will be among the casualties, but many of them will be.

**Staying Healthy:** Smoking cigarettes is risky, but it does offer several short-term benefits. They help you stay calm and relaxed in moments of stress, they keep you alert, and they help prevent overeating. Moreover, there is always someone's Aunt Edna who smoked four packs a day and lived to be 95. Besides, it is hard to kick the habit (over 70% of people who quit are back on cigarettes within three months). Are they so dangerous that it's worth trying to quit?

- Similar long-run risks: overeating or excessive drinking pose a risk over the long run. But a little alcohol each day is relaxing and may be good for your health.

- Similar short-term risks: drinking too much before driving or having sex without a condom pose a risk even in the short run.

**Raising your Children:** There are many gambles here. How strict should you be? What risks should you let them take? How much should you allow them to make their own decisions?

**Insurance:** When you buy insurance for your car or your home or yourself, the insurance company is betting that you will not need it. You are hedging your bets by buying it. (This is one bet you hope to *lose*).

**Starting a Business:** You would like the independence of owning your own business. Moreover, some small businesses do very well and you could make a lot of money. But many new businesses fail. Is it worth the risk? How can you even determine what the risk is?

**Drilling for Oil:** Drilling a new well is a risky proposition. It is expensive, and many wells never produce. But some locations are much more likely to yield oil than others. The relevant facts in determining such probabilities include the geological makeup of a region and the number of successful wells in the area.

**The Used Car Dealer:** You show up on the lot and want to buy a used pickup truck. It looks o.k., you hear that it only has 60,000 miles on it, and the price is right. But there isn't any long-term warranty.

**Seat Belts:** Seat belts save many lives. But some people die because they wore a seat belt. Moreover, if you don't like to be bothered putting them on, they involve costs in time and irritation. And even if you wear seatbelts, you run some risk every time you go somewhere in your car.

- There are many examples of this sort: wearing a helmet when you ride a bike or a motorcycle, driving way above the speed limit.

**Investments:** If you begin saving money soon after you graduate, you will have fewer worries about putting your children through college and you will provide for your retirement years. But as the collapse of Enron reminds us, the investments that promise the most gains are typically the most risky. What should you do?

**The Stock Market:** No one is sure what the market will do. If it does well, you can make much more money than you could with most other investments. But if the market takes a big dip, like it did in October of 1997 or April of 2000, you can also lose your shirt.

In short, we all have to make important decisions that involve uncertainty. For college students these include questions about what to major in, which courses to take, whether to go to graduate school, which courses to take, whether to study for the exam or to hit the bars, whether to break up with someone, and so on.

## 12.2 Inductively Strong Arguments

**Inductive strength:** if premises all true, a high probability the conclusion is true

When things are uncertain in these ways, we usually cannot expect to find deductively valid arguments. At most we can hope to find arguments that are inductively strong. In an earlier chapter we saw that an argument is **inductively strong** just in case:

1. If all of its premises are true, then there is a *high probability* that its conclusion will be true as well.
2. It is not deductively valid.

The first item is the important one (the point of the second item is to insure that no argument is both deductively valid *and* inductively strong; this makes things easier for us in various ways).

There are two important ways in which Inductive strength differs from deductive validity:

1. Unlike deductive validity, inductive strength *comes in degrees*.
2. In a deductively valid argument, the conclusion does not contain any information that was not already present in the premises. By contrast, in an inductively strong argument, the *conclusion contains new information*

Since the conclusion contains new information, we go beyond the information that is stated in our premises. Inductively strong arguments and reasons can take many different forms; in this module we will focus on those that involve probability.

We can also speak of *inductively strong reasons*. A group of sentences provide inductive reasons for a conclusion just in case it is unlikely for all of them to be true and the conclusion false. There is always an *inductive leap* from the inductively strong reasons to the conclusion. The stronger the inductive reasons, the less risky the inductive leap.

### Kinds of Reasoning

We can make a sharp distinction between deductively valid arguments, on the one hand, and those that are merely inductively strong, on the other, and it is important



to be clear about the difference. But in everyday life there is often no very clear distinction between deductive and inductive reasoning. What might seem to be invalid may become so if we supply plausible missing premises, for example. Still, it is clear that a great deal of our reasoning involves arguments and evidence that are inductively uncertain. Life is full of risks and uncertainty, and so is our reasoning about it. No methods are foolproof, but some are much better than others.

### 12.3 Chapter Exercises

1. List three real-life cases that involve probabilities and gambles. How do you try to determine how likely various outcomes are in these cases?
2. How What are some of the factors that are relevant in trying to decide whether to quit smoking?
3. Would you want to live near a nuclear power plant? How dangerous do you think such plants are? How could you find out more about how hazardous they are?
4. What connections are there between probabilities and the assessment of risks (like being in a automobile accident or receiving anthrax-contaminated mail)?
5. If an argument is deductively valid, then adding additional premises to it cannot destroy its validity. By contrast, inductively strong arguments can be weakened by adding the right sorts of premises. Given an example of how an argument that isn't valid but is inductively strong can be made weaker, then stronger, then weaker again by the addition of premises.



# Chapter 13

## Rules for Calculating Probabilities

**Overview:** In this chapter we introduce notation for expressing claims about probabilities and learn six rules for calculating the probabilities of three important types of sentences.

### Contents

---

<b>13.1 Intuitive Illustrations</b> . . . . .	<b>230</b>
<b>13.2 Probabilities are Numbers</b> . . . . .	<b>230</b>
13.2.1 Notation . . . . .	231
<b>13.3 Rules for Calculating Probabilities</b> . . . . .	<b>233</b>
13.3.1 Absolutely Certain Outcomes . . . . .	233
13.3.2 Negations . . . . .	234
13.3.3 Disjunctions with Incompatible Disjuncts . . . . .	236
<b>13.4 More Rules for Calculating Probabilities</b> . . . . .	<b>240</b>
13.4.1 Conjunctions with Independent Conjuncts . . . . .	240
13.4.2 Disjunctions with Compatible Disjuncts . . . . .	243
<b>13.5 Chapter Exercises</b> . . . . .	<b>244</b>
<b>13.6 Appendix: Working with Fractions</b> . . . . .	<b>246</b>

---

## 13.1 Intuitive Illustrations

### Drawing Balls from an Urn

Suppose we have an urn containing twenty balls. Each ball is either red or green (and no ball is both). Suppose that there are eight red balls and twelve green balls in the urn.

**Red Balls:** 8

**Green Balls:** 12

You mix the balls thoroughly, reach in and, without looking, draw one ball from the urn. What is the probability that you draw a red ball? There are twenty balls, eight of which are red. So there are eight ways of getting a red ball out of twenty possible cases. We can express this as a fraction,  $8/20$  (which reduces to  $2/5$ ). We will say that the probability of drawing a red ball is  $2/5$ .

What is the probability of drawing a green ball? Since there are twelve green balls out of the twenty, there are twelve chances in twenty. So the probability of drawing a green ball is  $12/20$ , which reduces to  $3/5$ .

What is the probability that you will draw *either* a red ball or else a green ball? Since all of the balls are either red or green, you are certain to pick one color or the other. The probability is  $20/20$ , which reduces to 1.

What is the probability that you will draw a ball that is neither red nor green? None of the balls satisfy this description. Put another way, 0 out of 20 balls satisfy it. So the probability is  $0/20$  which is just 0.

### Rolling Dice

You are rolling a fair die (one that isn't loaded). What are the chances that you will roll a 3? A die has six sides, marked with the numbers 1, 2, 3, 4, 5, and 6. So there is one chance out of six of getting a 3. We say that the probability of a 3 is  $1/6$ .

What are the chances of rolling an even number? There are three ways to roll an even number: rolling a 2, 4, or 6. So the probability is  $3/6$  (which reduces to  $1/2$ ).

## 13.2 Probabilities are Numbers

We think in terms of probabilities more than you might suppose. We often do this, for example, when we talk about percentages. When something could just as easily turn out either of two ways, we say that it's a fifty-fifty proposition. We

could express the same claim by saying that there is a probability of .5 that either of the two outcomes will occur. The meteorologist says that there is a 65% chance of showers later today. We could instead say that the probability of showers is .65.

In general, we can translate claims about percentages into corresponding claims about probability by dividing the percentage by 100; this just means adding a decimal point (and perhaps one or more zeros) at the appropriate place. Thus 90% means a probability of .9, and 3% means a probability of .03.

Polls and surveys also report percentages that can be translated into probabilities. Many of the sciences, from psychology to genetics to physics, make heavy use of probability and statistics (which is itself based on the theory of probability).

We are all constantly concerned with the likelihood of various possibilities many times everyday. Is rain *so* likely that we should cancel the hike? Are allergy shots sufficiently likely to help with my allergies that they are worth the time and bother? How likely is this prisoner to commit another crime if she is paroled this year? How likely is Wilbur to go out with Wilma if she asks him for a date?

Probabilities are numbers that represent the likelihood that something will happen. Probabilities are measured on a scale from 0 to 1. If something is certain to happen, it has a 100% chance of occurring, and we say that it has a probability of 1. And if something is certain *not* to happen, it has a 0% chance of occurring, and we say that it has a probability 0. The numbers 0 (“no way”) and 1 (“for sure”) nail down the end points of this scale, so it is impossible to have probability values of 1.3,  $-2$ , or  $5/4$ . Since life is usually uncertain, we are often dealing with probabilities greater than 0 and less than 1.

Probabilities are numbers from 0 to 1

Probabilities measure how likely things are

### 13.2.1 Notation

In our first example, the probability of drawing a red ball is  $2/5$ . We might write this as: “Probability(Drawing a red ball) =  $2/5$ .” But it will save a lot of writing if we introduce two sorts of abbreviations. First, we abbreviate the word ‘probability’ with ‘Pr’. Second, we abbreviate sentences by capital letters. You can use any letters you please (so long as you do not use one letter to abbreviate two different sentences in a problem). But it’s best to pick a letter that helps you remember the original sentence. For example, it would be natural to abbreviate the sentence ‘I rolled a six’ by ‘S’.

If we abbreviate the sentence ‘I drew a red ball’ as ‘R’, we can write our claim that the probability of drawing a red ball is  $2/5$  like this:

$$\text{Pr}(R) = \frac{2}{5}$$

If we flip a fair coin it is equally likely to come up heads or tails, so we say that

the probability of getting heads on the next toss is .5. We write this as:  $\Pr(H) = .5$  (or  $\Pr(H) = 1/2$ ). In general, we write

$$\Pr(S) = n$$

to mean that the sentence  $S$  has a probability of  $n$  of being true.<sup>1</sup>

### Exercises

How do we use this notation to express the following claims:

1. The probability of rolling a six is  $1/6$ ?
2. The probability of drawing an ace of spaces from a full deck is  $1/52$ ?
3. The probability of rolling a five *or* a six is  $2/6$  (i.e.,  $1/3$ )?
4. The probability of drawing either a red ball or a green ball is 1?
5. The probability of getting both a red ball and a green ball on the same draw is 0?

### Answers

There isn't a uniquely correct way to abbreviate the simple sentences in this exercise, but the following ways are natural.

1.  $\Pr(S) = 1/6$
2.  $\Pr(A_s) = 1/52$
3. We have to sneak an 'or' into our abbreviated sentence:  $\Pr(F \text{ or } S) = 1/3$
4.  $\Pr(R \text{ or } G) = 1$
5.  $\Pr(R \text{ and } G) = 0$

### Cards and Dice: The Basics

- ♠ ♣ Some of the problems we will consider involve cards and dice; here is the makeup of a standard deck of cards (with the jokers removed) and the possible outcomes when you roll a pair of dice.

<sup>1</sup>Sometimes it is most natural to speak of the probability of a *sentence* being true or false. Other times it is more natural to speak of the probability of a given *event* occurring or some process having a particular outcome. In a more advanced treatment these things would matter, and there we could construe talk about the occurrence of events as shorthand for claims that particular sentences describing those events are true. But we don't need to worry about niceties like these in this course.

♣	◇	♥	♠
Ace	Ace	Ace	Ace
King	King	King	King
Queen	Queen	Queen	Queen
Jack	Jack	Jack	Jack
10	10	10	10
9	9	9	9
8	8	8	8
7	7	7	7
6	6	6	6
5	5	5	5
4	4	4	4
3	3	3	3
2	2	2	2

There are fifty two cards, thirteen in each suit (with the jokers removed).

Figure 13.1: Makeup of a Standard Deck of Cards

Die One:	1	2	3	4	5	6
Die Two:	1	2	3	4	5	6

Figure 13.2: Outcomes of Rolling Dice

### 13.3 Rules for Calculating Probabilities

Because probabilities are numbers, we have to use a bit arithmetic to calculate them. Don't worry if numbers make you nervous; we will only need things like multiplying fractions which you learned long ago. Still, it may have been awhile since you worked with fractions, so if you don't feel confident about them take a few minutes to work through the appendix (p. 246), which reviews the basic arithmetic that you'll need.

#### 13.3.1 Absolutely Certain Outcomes

We now introduce eight rules that will help us calculate probabilities. It is important that you learn and understand these rules. If you don't, you simply won't be

able to work the problems.<sup>2</sup>

**Rule 1. (for events that are certain to occur):** If something is certain to happen, its probability is 1. If the sentence  $A$  is certain to be true:

$$\Pr(A) = 1$$

*Example:* If you draw a ball out of the urn described above, you are certain to get either a red ball or a green ball:  $\Pr(R \text{ or } G) = 1$ .

**Rule 2. (for events that are certain *not* to occur):** If something is certain *not* to happen, its probability is 0. If the sentence  $A$  is certain to be false:

$$\Pr(A) = 0$$

*Example:* If you draw a ball out of the urn, there is *no way* that you will get a ball that is *both red and green*;  $\Pr(R \ \& \ G) = 0$ .

### 13.3.2 Negations

The negation of a sentence says that the negated sentence is false. For example, ‘I did not draw a red ball’ negates the sentence ‘I drew a red ball’. We will use  $\neg$  to signify negation. So we express the negation of the sentence  $S$  by writing  $\neg S$ .

*Example:* If ‘ $A$ ’ stands for the claim that I drew an ace,  $\neg A$  says that I did *not* draw an ace.

Probabilities and of sentences and their negations are like people on a seesaw (Figure 13.3 on the next page). The lower you go, the higher the person on the other side goes. And the higher they go, the lower you go. Similarly, the lower the probability of a sentence, the higher the probability of its negation. And the higher the probability of a sentence, the lower the probability of its negation. If you come to think it more likely that you will pass Chemistry 101 you should think it less likely that you will fail.

The “amount” of probability is limited. A sentence and its negation have a total probability of 1 to divide between them. So whatever portion of this doesn’t go to

<sup>2</sup>The presentation of the rules for probability follows that of Brian Skyrms’s excellent book *Choice & Chance: An Introduction to Inductive Logic*, Wadsworth, 4th/ed, 2000. The muddy Venn diagrams are from Bas C. van Fraassen’s *Laws and Symmetry*, Oxford: Clarendon Press, 1989.






Figure 13.3: A Sentence and its Negation Split One Unit of Probability

a sentence goes to its negation. In other words, the probabilities  $S$  and  $\neg S$  always add up to 1.

**Rule 3. (negations):** The probability of a negation is 1 minus the probability of the negated sentence.

$$\Pr(\neg A) = 1 - \Pr(A)$$

*Example 1:* If the probability of drawing an ace is  $1/13$ , then the probability that you will *not* draw an ace is  $12/13$ .

*Example 2:* If a coin is bent so that the probability of tossing heads is  $.4$ , then the probability of not getting a head on a toss is  $.6$ .

The circle labeled  $A$  in Figure 13.4 on the following page represents the cases in which  $A$  is true. For example, it might mean that we draw an ace from a deck of cards. The region of the rectangle that is not in  $A$  represents the negation of  $A$ . The rectangle represents a total probability of 1, and it represents the amount of the rectangle not in  $A$  is 1 minus the amount of the rectangle that is in  $A$ .

In simple cases we can represent probabilities by Venn diagrams like that in Figure 13.4. The rectangle represents all the things that could possibly happen. It has a total probability of 1. Think of it as having one bucket, one unit, of mud spread over its surface. The mud represents the probability. Several situations are possible in Figure 13.4.

- All of the mud might be *inside* the circle  $A$ ; this represents the case where the probability of  $A$  is 1 (it has the entire unit) and that of  $\neg A$  is 0.
- All the mud might be *outside* the circle  $A$ ; this represents the case where the probability of  $A$  is 0 and that of  $\neg A$  is 1 (it has the entire unit).

- Some mud may be inside  $A$  and some outside. Then neither  $A$  nor  $\neg A$  have probabilities of 1 or of 0. The more mud inside  $A$ , the more probable it is.




Figure 13.4: Negations

### Exercises

1. Suppose  $2/3$  of the mud in Figure 13.4 is placed inside circle  $A$ . What are the probabilities of  $A$  and  $\neg A$  given this representation?
2. Suppose virtually all of the mud in Figure 13.4 is placed outside circle  $A$ . What does this tell us about the relationship between the probabilities of  $A$  and  $\neg A$ ?

### 13.3.3 Disjunctions with Incompatible Disjuncts

A **disjunction** is an “either/or” sentence. It claims that either, *or both*, of two alternatives is the case. Here are two specimens:

1. Either the butler did it or the witness for the defense is lying.
2. Either I’ll roll a five or I’ll roll a six.

The two simpler sentences that make up a disjunction are called *disjuncts*. The order of the disjuncts in a disjunction doesn’t matter. Note that we interpret disjunctions so that they are true if both disjuncts are true. Our *or* has the same meaning as the phrase *and/or*, so a disjunction claims that *at least one* of the disjuncts is true.

### Incompatibility

Two things are **incompatible** just in case they cannot both occur (or cannot both be true) together. It is impossible for them both to happen in any given situation. The

truth of either *excludes* the truth of the other, so incompatible things are sometimes said to be mutually exclusive.

Incompatibility is a two-way street: if one thing is incompatible with a second, the second is incompatible with the first. If  $A$  and  $B$  are incompatible, then *no*  $A$ s are  $B$ s, and no  $B$ s are  $A$ s. So if  $A$  and  $B$  are incompatible,  $\Pr(A \& B) = 0$ .

*Example 1:* Getting a head on the next toss of a coin and getting a tail on that *same* toss are incompatible. Getting either excludes getting the other.

*Example 2:* Getting a head on this toss and getting a tail on the *subsequent* toss are *compatible*. These two outcomes are in no way inconsistent with each other. Neither precludes the other.

**Incompatible sentences:**  
cannot both be true at the same time

### Exercises

Which of the following pairs are incompatible with each other?

1. Getting a 1 on the next roll of die. Getting a 3 on that same roll.
2. Getting a 1 on the next roll of die. Getting a 3 on the roll after that.
3. Wilbur graduates from OU this spring. Wilbur fulfills his life-long dream and begins a career as a movie usher.
4. Wilbur graduates from OU this spring. Wilbur flunks out of OU this spring.
5. Wilbur turns twenty. On that very day he gets the good news that he has just become the President of the United States.
6. Wilbur passes all of the exams in this course. Wilbur passes the course.
7. Wilbur gets a very low  $F$  on all of the exams in this course. Wilbur passes the course.

### Answers

1. Incompatible. You can't get a 1 *and* a 3 on the very same roll.
2. Compatible. No side of the die has both a 1 and a 3 on it.
3. Compatible.
4. Incompatible. Graduating and flunking out exclude each other; if either happens, the other cannot.
5. Incompatible. The President has to be at least thirty-five. So being twenty and being President preclude each other. You can't be both at once.
6. Compatible
7. What do you think?

### The Probability of a Disjunction with Incompatible Disjuncts

What is the probability that a disjunction,  $A$  or  $B$ , with *incompatible* disjuncts is true? We can represent the situation with Figure 13.5.




Figure 13.5: Disjunctions with Incompatible Disjuncts

Our question about the probability of the disjunction  $A$  or  $B$  now translates into the question: What is the total area occupied by the two circles? And the answer is: it is just the area occupied by  $A$  added to the area occupied by  $B$ . In terms of muddy diagrams, we take the total amount of mud that is on either  $A$  or on  $B$  and add them together.

**Rule 4. (disjunctions with incompatible disjuncts):** The probability that any disjunction with *incompatible* disjuncts is true is the sum of the probabilities of the two disjuncts.

$$\Pr(A \text{ or } B) = \Pr(A) + \Pr(B).$$

*Example:* No card in a standard deck is both an ace and a jack. So drawing an ace is incompatible with drawing a jack. If the probability of drawing an ace is  $1/13$  and the probability of drawing a jack is  $1/13$ , then the probability of drawing either an ace or a jack is  $1/13 + 1/13 = 2/13$ .

We can extend our rule to disjunctions with more than two alternatives (disjuncts). As long as each disjunct is incompatible with all of the other disjuncts, we can determine the probability of the entire disjunction by adding the individual probabilities of each of its disjuncts. For example, the probability that I will draw either a king or a queen or a jack on a given draw is  $1/13 + 1/13 + 1/13 (= 3/13)$ .

### Exercises

1. Remove the jokers from a standard deck of playing cards, so that you have 52 cards. You are drawing one card at a time (and each card has an equally good

chance of being drawn). What is the probability of drawing each of the following? In cases where more than a single card is involved, specify which rules are relevant (you will be able to calculate some of these without using the rules, but you won't be able to do that when we get to harder problems, so it is important to begin using the rules now).

1. A jack of diamonds
  2. A jack
  3. A king or a jack
  4. A two of clubs
  5. The jack of diamonds or the two of clubs
  6. A red jack
  7. A card that is not a red jack
  8. A face card (king, queen or jack) or an ace
  9. A card that is either a face card or else not a face card
  10. A card that is both a face card and not a face card
2. You are going to roll a single die. What is the probability of throwing:
1. A one
  2. A three
  3. A one or a three
  4. An even number
  5. A non-three
  6. A two or a non-even number
3. With Rule 4 (our new rule for disjunctions) and Rule 1 (our rule for sentences that have to be true) we can prove that R3, our rule for negations, is correct. Try it.

### Answers to Selected Exercises

3. Here's how to use Rule 4 and Rule 1 to show that R3, our rule for negations, is correct. First note that each sentence is incompatible with its negation, so  $A$  and  $\neg A$  are incompatible. Moreover, the sentence ' $A$  or  $\neg A$ ' is certain to be true. Hence

1.  $\Pr(\neg A \text{ or } A) = 1$  [by Rule 1]
2.  $\Pr(\neg A \text{ or } A) = \Pr(A) + \Pr(\neg A)$  [by Rule 4]
3. So  $\Pr(\neg A) + \Pr(A) = 1$  [from 1 and 2]
4. Hence,  $\Pr(\neg A) = 1 - \Pr(A)$  [by subtracting  $\Pr(A)$  from both sides]

## 13.4 More Rules for Calculating Probabilities

### 13.4.1 Conjunctions with Independent Conjuncts

A **conjunction** is an *and*-sentence. The sentence ‘Wilbur passed the final and Betty passed the final’ is a conjunction. The two simpler sentences glued together by the ‘and’ are called *conjuncts* (the order of conjuncts in a conjunction doesn’t matter). A conjunction is true just in case *both* of its conjuncts are true; if either conjunct is false, the whole thing is false. We will use ‘&’ to abbreviate ‘and’.

#### Independence

**Independent sentences:** are completely unconnected, irrelevant to one another

Two sentences are **independent** (of each other) just in case they are completely irrelevant to each other. The truth-value of one has no effect or influence or bearing on the truth value of the other. Knowing that one is true (or false) tells you nothing about whether the other is true (or false). Independence is a two-way street: if one thing is independent of a second, the second is also independent of the first.

*Example 1:* You are drawing cards from a deck, and after each draw you replace the card and reshuffle the deck. The results of the two draws are independent. What you get on the first draw has no influence on what you get on the second.

*Example 2:* You are drawing cards from a deck without replacing them. What you get on the first draw changes the makeup of the deck, and so the outcome of the first draw does bear on the outcome of the second. The outcome of the second draw is (to some degree) *dependent* on the outcome of the first.

Do not confuse incompatibility with independence. They are completely different.

1. Two things are **incompatible** just in case they cannot both be true at the same time; the truth of either *excludes* the truth of the other.
2. Two things are **independent** just in case the truth value of each has no bearing on the truth value of the other.

Example: Getting a head on the next toss of a coin and a tail on that *same* toss are incompatible. But they are *not* independent.

#### Exercises

Which of the following pairs are *incompatible*? Which are *independent*?

1. Getting a 1 on the next roll of a die. Getting a 3 on that same roll.
2. Drawing an ace on the first draw from a deck. Drawing a Jack on that same draw.

3. Getting a 1 on the next roll of a die. Getting a 3 on the roll after that.
4. Ross Perot is elected President in the next election. I roll a 3 on the first roll of a die.
5. Getting a head on the next flip on a coin. Getting head on the flip after that.
6. Passing all of the exams in this course. Passing the course itself.

### Rule for Conjunctions with Independent Conjuncts

**Rule 5. (conjunctions with independent conjuncts):** If the sentences  $A$  and  $B$  are independent, then the probability that their conjunction,  $A \& B$ , is true, is  $\Pr(A)$  times  $\Pr(B)$ .

$$\Pr(A \& B) = \Pr(A) \times \Pr(B)$$

So when two things are independent, the probability of their joint occurrence is determined by the simple multiplicative rule: multiply the probability of one by the probability of the other.

*Example:* What happens on the first toss of a coin has no effect on what happens on the second; getting a head on the first toss of a coin ( $H_1$ ) and getting a head on the second toss ( $H_2$ ) are independent. Hence,  $\Pr(H_1 \& H_2) = \Pr(H_1) \times \Pr(H_2) = 1/2 \times 1/2 (= 1/4)$




Figure 13.6: Tree Representation of the Probability of a Conjunction

The tree diagram (Figure 13.6) represents the possible outcomes. The numbers along each path represent the probabilities. The probability of a heads on the first flip (represented by the first node of the top path) is  $1/2$ , and the probability of a second head (represented by the node at the upper right) is also  $1/2$ . There are four paths through the tree, and each represents one possible outcome. Since all of the

four paths are equally likely, the probability of going down any particular one is  $1/4$ .

We can present the same information in a table (Figure 13.7) that shows more clearly why we *multiply* the probabilities of the two conjuncts. The outcomes along the side represent the two possible outcomes on the first toss, and the outcomes along the top the two outcomes of the second toss.

	$H_2$	$T_2$
$H_1$	$H_1 \& H_2$	$H_1 \& T_2$
$T_1$	$T_1 \& H_2$	$T_1 \& T_2$

Figure 13.7: Table Representation of the Probability of a Conjunction

We can extend our rule to conjunctions with more than two conjuncts. As long as each conjunct is independent of all the rest, we can determine the probability of the entire conjunction by multiplying the individual probabilities of each of its conjuncts. For example, the probability that I will get heads on three successive flips of a coin is  $1/2 \times 1/2 \times 1/2$ .

Our work will be much simpler because of the following facts.

**Incompatibility:** only matters for disjunctions

- **Incompatibility is only relevant for disjunctions.** We do not need to worry about whether the conjuncts of a conjunction are incompatible or not.

**Independence:** only matters for conjunctions

- **Independence is only relevant for conjunctions.** We do not need to worry about whether the disjuncts of a disjunction are independent or not.

### Winning the Lottery

Rich folks put money in a retirement fund. Rednecks play the lottery.  
— Jeff Foxworthy

The chances of winning a state lottery are very low; you have much better chances of winning in almost any casino in the world. To see why, imagine a lottery where you have to correctly guess a 1-digit number. There are 10 such digits, so your chances are 1 in 10, or .1. So far, so good. But now imagine that you must guess a two-digit number. There are ten possibilities for the first digit *and* ten possibilities of the second. Assuming the two digits are independent, this means that the chances of correctly guessing the first digit *and* the second digit are  $\frac{1}{10} \times \frac{1}{10} = \frac{1}{100}$ . You would win this lottery about once every 100 times you played.

This may not sound so bad. But most state lotteries require you to match about *twelve* one-digit numbers. In this case, we determine the probability of winning by



multiplying  $\frac{1}{10}$  by itself *twelve times*. When we write  $\frac{1}{10^{12}}$  out the long way, it turns out to be

$$\frac{1}{1,000,000,000,000},$$

which is almost infinitesimally small.

### 13.4.2 Disjunctions with Compatible Disjuncts

Whenever the disjuncts of a disjunction are *incompatible*, R4 applies, but when they are compatible, we need a more subtle rule.

It will help you to see why if we consider the following example. We are going to flip a quarter twice. What is the probability of getting a head on at least one of the two tosses; what is  $\Pr(H_1 \text{ or } H_2)$ ? The probability of getting heads on any particular toss is  $1/2$ . So if we used our old disjunction rule (Rule 4., for incompatible disjuncts), we would have  $\Pr(H_1 \text{ or } H_2) = \Pr(H_1) + \Pr(H_2)$ , which is just  $1/2 + 1/2$ , or 1. This would mean that we were *certain* to get a head on at least one of our two tosses. But this is obviously incorrect since it is quite possible to get two tails in a row.

Indeed, if we used our old disjunction rule to calculate the probability of getting a head on at least one of three tosses, we would have  $1/2 + 1/2 + 1/2$ , which would give us a probability greater than 1.5 (and this could never be correct, since probabilities can *never* be greater than 1).




Figure 13.8: Disjunctions with Compatible (“Overlapping”) Disjuncts

If  $A$  and  $B$  are compatible, it is possible that they could occur together. For example, drawing an ace and drawing a black card are compatible (we might draw the ace of spades or the ace of clubs). We indicate this in Figure 13.8 by making the circle representing  $A$  and the circle representing  $B$  overlap. The overlapping, cross-hatched, region represents the cases where  $A$  and  $B$  overlap.

In terms of muddy diagrams we add the weight of the mud on  $A$  to the weight of the mud on  $B$ , but when we do this we weigh the mud where they overlap twice.

So we must subtract once to undo this double counting. We must subtract the probability that  $A$  and  $B$  both occur so that this area only gets counted once.

### The General Disjunction Rule

**Rule 6. (disjunctions):** The probability of any disjunction, incompatible or compatible, is the *sum* of the probabilities of the two disjuncts, *minus* the probability that they both occur.

$$\Pr(A \text{ or } B) = \Pr(A) + \Pr(B) - \Pr(A \& B)$$

*Example 1:* Drawing an ace and drawing a club are not incompatible. So  $\Pr(A \text{ or } C) = \Pr(A) + \Pr(C) - \Pr(A \& C)$ ; so it equals  $1/13 + 1/4 - 1/52$ . We subtract the  $1/52$  because otherwise we would be counting the ace of clubs twice (once when we counted the aces, and a second time when we counted the clubs).

*Example 2:* Getting heads on the first and second flips of a coin are compatible. So to calculate  $\Pr(H_1 \text{ or } H_2)$  we have to subtract the probability that both conjuncts are true. We must consider  $\Pr(H_1) + \Pr(H_2) - \Pr(H_1 \& H_2)$ , which is  $1/2 + 1/2 - 1/4 (= 3/4)$ .

Rule 6 is completely general; it applies to *all* disjunctions. But when the two disjuncts are *incompatible*, the probability that they are both true is 0, so we can forget about subtracting anything out.

## 13.5 Chapter Exercises

1. You roll a pair of dice. Assume that the number that comes up on each die is independent of the number what comes up on the other (which is the case in all normal situations).
  1. What is the probability that you roll two sixes (“box cars”). Hint: this is a “one-way” point; both dies must come up sixes.
  2. What is the probability that you roll two ones?
  3. What is the probability that you not roll a double six?
  4. What is the probability that you will either roll two sixes or else roll two ones?
  5. What is the probability that you roll a five?
  6. What is the probability that you roll a seven or eleven?

2. You are going to draw one card from a standard deck of playing cards. Once you see what the card is, you replace it, then draw a second card. Determine the probabilities of each of the following, say which rules are relevant, and explain how you use the rules to obtain the results.

1. What is the probability that you get a jack on the first draw?
2. What is the probability that you get a diamond on the first draw?
3. What is the probability that you get a jack of diamonds on the first draw?
4. What is the probability that you get a jack *or* a diamond on the first draw?
5. What is the probability of a queen on the first draw and a jack on the second?
6. What is the probability of getting one jack and one queen (here the order in which you get them doesn't matter)?
7. What is the probability of getting two aces?
8. What is the probability of drawing exactly one ace?
9. What is the probability of getting at least one ace?
10. What is the probability of not getting an ace on either draw?

3.  $\Pr(P) = 1/2$ ,  $\Pr(Q) = 1/2$ , and  $\Pr(P \& Q) = 1/4$ .

1. Are  $P$  and  $Q$  incompatible? Why or why not?
2. What is the probability of  $P$  or  $Q$ ?

4. Consider the following possible outcomes of flipping a coin three times, where H = Head and T = Tail.

HHH, TTT, HTH, THT

You know that if the coin is fair the probability of all four sets of outcomes is the same:  $1/2 \times 1/2 \times 1/2 = 1/8$ . Now calculate the probabilities for each of the outcomes when the coin is biased and the probability of getting a head on any flip is 0.70.

5. What is the probability of getting a head on the first flip or on the second when you flip the biased coin described in the previous problem two times?

## 13.6 Appendix: Working with Fractions

### But I hate math ...

You knew all the arithmetic you will need for this course by the end of ninth grade, but it's easy to get rusty. Don't worry if you are, but do review the following material. If you have a little "math anxiety," keep in mind that the key is to approach things *slowly*. Each of the basic concepts is relatively easy, and if you work to understand each point before going on to the next, you will be able to master the material. In fact, the only algebra you will need is the very minimal amount required to add and multiply rather simple fractions.

Try to think of this in terms of a number of small steps, rather than trying to grasp everything all at one. As with much else in this course, you will also need to work a number of problems on your own. The most important thing in mastering any skill is practice.

### How Fractions Work

A fraction consists of a *numerator* and a *denominator*. The numerator is the number on the top and the denominator is the number on the bottom. So the numerator of  $5/7$  is 5, and the denominator is 7. Two fractions have a common denominator just in case they have the same denominator;  $5/7$  and  $3/7$  have a common denominator (namely 7), but  $5/7$  and  $8/11$  do not.

It is often easier to work with fractions if we convert them into their decimal equivalents. To find the decimal equivalent for a fraction, divide the numerator by the denominator. For example, to convert  $1/4$  to a decimal, divide 1 by 4 (to get .25). To convert  $3/5$  to a decimal, divide 3 by 5 (to get .6). Such conversions are easy if you use a calculator (which you are encouraged to do).

### Adding Fractions

To *add* two fractions that have a common denominator, you simply add their numerators and write it above their denominator. For example,  $3/7 + 2/7 = 5/7$ . And  $4/52 + 3/52 = 7/52$ .

If you want to add fractions that have different denominators, you must find a common denominator. Once you do this, you simply add their numerators and write the result above the common denominator. In many cases finding a common denominator is straightforward, but you can avoid such worries if you replace the fractions by their decimal equivalents and simply add those.

*Example:* Add  $3/5 + 1/4$ . You can either find a common denominator or you can add their decimal equivalents.

**Common Denominator** The lowest common denominator of  $3/5$  and  $1/4$  is 20. So we can express  $3/5$  as  $12/20$  and  $1/4$  as  $5/20$ . And  $12/20 + 4/20 = 17/20$ .

**Decimal Equivalents** The decimal equivalent of  $3/5$  is .6 (divide 3 by 6 to get this) and the decimal equivalent of  $1/4$  is .25 (divide 1 by 4). So  $3/5 + 2/3 = .6 + .25 = .85$

We can check our two approaches by seeing whether they yield the same result; is .85 equal to  $17/20$ ? To answer this we divide 17 by 20, which is .85, just as it should be.

### Multiplying fractions

To *multiply* fractions you just multiply their numerators to get the new numerator and you multiply their denominators to get the new denominator.

#### Examples

1. What is  $3/5 \times 3/4$ ? Multiply the two numerators ( $3 \times 3$ ) to get the new numerator, which is 9, then multiply the two denominators ( $5 \times 4$ ) to get the new denominator, which is 20. Putting these together, the answer is  $9/20$ .
2. What is  $4/52 \times 3/51$ ? Multiply the numerators to get 12 and the denominators to get 2652. So the answer is  $12/2652$  (which reduces to  $1/222$ ).

You can also always convert fractions to their decimal equivalents and then multiply them. We won't worry much at the beginning about reducing fractions. But do note that when you multiply fractions you must multiply their denominators as well as their numerators.  $4/52 \times 3/52$  is *not*  $12/52$  (it's  $12/(52 \times 52)$ ).

Probabilities range from zero to one, and most of our calculations will involve fractions between zero and one. There are two very important points to remember about such fractions.

1. When you *add* one such fraction to another, the result will be *larger* than either fraction alone.
2. When you *multiply* one such fraction by another, the result will be *smaller* than either fraction alone.

**Exercises**

Find the value of each of the following:

1.  $\frac{2}{3} + \frac{1}{3}$
2.  $\frac{2}{6} + \frac{1}{6}$
3.  $\frac{2}{3} + \frac{1}{6}$  (you need a common denominator here )
4.  $\frac{4}{9} + \frac{11}{20}$
5.  $\frac{2}{3} \times \frac{1}{3}$  (the denominator here will be 9, *not* 3.
6.  $\frac{2}{6} \times \frac{1}{6}$
7.  $\frac{2}{3} \times \frac{1}{6}$  (when we multiply fractions we *don't* use a common denominator).
8.  $\frac{2}{6} \times \frac{1}{3}$
9.  $\frac{4}{9} \times \frac{11}{20}$
10.  $\frac{4}{52} \times \frac{1}{51}$

# Chapter 14

## Conditional Probabilities

**Overview:** In this final chapter on probability we will learn how to deal with conditional probabilities, the probabilities of conjunctions whose conjuncts are not independent, and the relationship between probabilities and odds.

### Contents

---

<b>14.1 Conditional Probabilities</b> . . . . .	<b>249</b>
14.1.1 Characterization of Conditional Probability . . . . .	251
14.1.2 The General Conjunction Rule . . . . .	253
<b>14.2 Analyzing Probability Problems</b> . . . . .	<b>254</b>
14.2.1 Examples of Problem Analysis . . . . .	255
<b>14.3 Odds and Ends</b> . . . . .	<b>259</b>
14.3.1 Sample Problems with Answers . . . . .	262
14.3.2 More Complex Problems . . . . .	263
<b>14.4 Chapter Exercises</b> . . . . .	<b>264</b>
14.4.1 Summary of Rules . . . . .	267

---

### 14.1 Conditional Probabilities

As the world changes, probabilities change too. The probability of drawing an ace from a full deck of cards is  $4/52$ . But if you draw two aces and don't replace them,

the probability of drawing an ace changes. We say that the *conditional probability* of drawing an ace *given* that two aces have been removed is  $2/50$ .

The probability of something being the case *given that* something else is the case is called a **conditional probability**. We express the conditional probability of  $A$  on  $B$  by writing  $\Pr(A|B)$ . We read this as ‘the probability of  $A$  *given*  $B$ ’. In the example above we are interested in the probability of drawing an ace *given that* two aces have already been drawn.

Much *learning* involves conditionalization. As we acquire new information, our assessments of probabilities change. You always thought Wilbur was very honest, but now you learn that he stole someone’s wallet and then lied about it. This leads you to reassess your belief that he has probably been honest on other occasions. You *conditionalize* on the new information about Wilbur, *updating* your views about how probable things are in light of the new evidence.

*Example 1:* Your friend asks you to pick a card, any card, from a full deck. How likely is it that you drew a king? Now your friend looks at the card and declares that it’s a face card. This new information changes your estimate of the probability that you picked a king. You are now concerned with the probability that you drew a king *given that* you drew a face card.

*Example 2:* The probability of getting lung cancer ( $C$ ) is higher for smokers ( $S$ ) than for nonsmokers. In our new notation this means that  $\Pr(C|S)$  is greater than  $\Pr(C|\neg S)$ .

*Example 3:* You are about to roll a fair die. The probability that you will roll a four is  $1/6$ . Your roll too hard and it tumbles off the table where you can’t see it, but Wilbur looks and announces that you rolled an even number. This thins the set of relevant outcomes by eliminating the three odd numbers. Figure 14.1 depicts the possibilities before and after Wilbur’s announcement. Before the announce-




Figure 14.1: Thinning the Relevant Outcomes

ment the probability of rolling a four was  $1/6$ . But once you thin out the relevant outcomes (by conditionalization), there are only three possibilities left, and only one way out of those three of rolling a four. When we restrict our attention in this way, now focusing only on the even numbers, we are said to *conditionalize* on the claim that the number is even.



### 14.1.1 Characterization of Conditional Probability

The next rule gives the definition for conditional probabilities.

**Rule 7. (conditional probability):** The probability of  $A$  given  $B$  is the probability of the conjunction of  $A$  &  $B$ , divided by the probability of  $B$ .

$$\Pr(A|B) = \frac{\Pr(A \& B)}{\Pr(B)}$$

In Rule 7 we must also require that the probability of  $B$  is *not zero* (because division by zero is undefined).

The idea behind Rule 7 is that conditional probabilities change the set of relevant outcomes. When your friend tells you that you selected a face card, the set of relevant possibilities shrinks from 52 (it might be *any* of the cards in the deck) down to 12 (we now know that it is one of the twelve face cards).

We put  $A \& B$  in the *numerator*, because we have now restricted the range of relevant cases to those covered by  $B$ . This means that the only relevant part of the region for  $A$  is the part that overlaps  $B$ , which is just the part where the conjunction  $A \& B$  is true. So in terms our diagrams,  $\Pr(A|B)$  is the amount of  $B$  occupied by  $A$ .

And we put  $\Pr(B)$  in the *denominator* because we want to restrict the range of relevant possibilities to those in which  $B$  is true. This is just what it means to talk about the probability of  $A$  *given*  $B$ . It may not be obvious that these numbers do the desired job, though, so we'll work through an example to see exactly how things work.

**How the Numbers Work** Suppose there are 100 students in your English class.<sup>1</sup> There are 50 men ( $M$ ), and 20 of them are Texans ( $T$ ). We can use these probabilities and Rule 7 to determine the probability of someone being a Texan given that they are male, i.e.,  $\Pr(T|M)$ . We have:

- $\Pr(T \& M) = \frac{20}{100}$  (the probability—or proportion—of people in the class who are male *and* Texans)
- $\Pr(M) = \frac{50}{100}$  (the probability—or proportion—of males in the class)

<sup>1</sup>This subsection is more difficult than most of the text and you can skip to §14.1.2 on the general conjunction rule without loss of continuity. You will understand the material better if you don't, though.

We then plug these numbers into the formula given by Rule 7 on the left to get the actual values at the right (Figure 14.2).




Figure 14.2: Conditionalization Trims out a New Unit

$$\underbrace{\Pr(T|M)}_{\text{By Rule 7}} = \underbrace{\frac{\Pr(T \& M)}{\Pr(M)}}_{\text{plugging in the numbers}} = \frac{\frac{20}{100}}{\frac{50}{100}} = \frac{20}{100} \times \frac{100}{50}$$

division by  $\frac{50}{100}$  = multiplication by  $\frac{100}{50}$

So the probability of someone in the class being a Texan if they are male is  $\frac{20}{100} \times \frac{100}{50} = \frac{20}{50}$  (the two 100s cancel)  $\frac{2}{5} = .4$ .

**What the Numerator Does** We disregard everyone who is not a male (some of whom may, but need not, be Texans). Figure 14.2 represents this by cutting out the circle of Males. We are then only interested in the percentage of Texans among males, which is given by the probability of someone in the class being both Texan and Male. We represent this as  $\Pr(T \& M)$ . It's just the overlap between the Texans and Males.

**What the Denominator Does**  $M$  only had half of the probability before, but once we focus on Males, once we conditionalize on this, trimming away everything else,

the probability of  $M$  should become 1. So we need to increase the probability of  $M$  from  $1/2$  (what it was before) to 1 (what it is once we confine attention to males).

Dividing by a fraction yields the same result as inverting and multiplying by it. So things work out because dividing by  $50/100$  is the same as multiply by  $100/50$ , i.e., it's the same as multiplying by 2. This ensures that we can treat  $M$  as now having the entire unit of probability (once we conditionalize on  $M$ ).

In terms of mud, when we shear off everything outside  $M$  we must also throw away all of the mud that was originally outside  $M$ . We then think of  $M$  as the new *total* area, and so we now view the amount of mud on *it* as *one unit*. Another way to see that  $M$  should now have a probability of 1 once we conditionalize on  $M$  is to note that  $\Pr(M|M) = 1$ .

In

$$\Pr(A|B) = \frac{\Pr(A \& B)}{\Pr(B)}$$

the less probable  $B$  was before we conditionalized, the more we have to multiple  $\Pr(A \& B)$  to inflate the new probability of  $B$  up to 1. If the probability of  $B$  was  $1/2$  we divide by  $1/2$ , which has the effect of multiplying by 2. If the probability of  $B$  was  $1/5$  we divide by  $1/5$ , which has the effect of multiplying by 5. Here  $1/5 \times 5/1$  gets us back to 1 unit of probability. In short division by the old  $\Pr(B)$  makes the new (post conditionalization)  $\Pr(B) = 1$ .

In general  $\Pr(A|B)$  is *not* equal to  $\Pr(B|A)$ . The probability that someone is a male given that he plays for the New York Yankees is 1. But the probability that someone is a Yankee given that he is male is very small. We will see in a later chapter that  $\Pr(A|B) = \Pr(B|A)$  just in case  $\Pr(A) = \Pr(B)$ . More importantly, we will see that confusing these two probabilities is responsible for a good deal of bad reasoning.

### 14.1.2 The General Conjunction Rule

By rearranging the terms in Rule 7, we obtain a general rule for conjunctions (divide both sides of the equality in Rule 7 by  $\Pr(B)$ ).

**Rule 8. (conjunctions):** The probability of the conjunction  $A \& B$ , where the conjuncts need not be independent, is the probability of  $A$  multiplied by the probability of  $B$  given  $A$ .

$$\Pr(A \& B) = \Pr(A) \times \Pr(B|A)$$

This rule is more general than Rule 5. It applies to *all* conjunctions, whether

their conjuncts are independent or not. Unlike Rule 7, we will often use Rule 8. in our calculations.

*Example:* You draw two cards from a full deck, and you don't replace the first card before drawing the second. The probability of getting a king on both of your draws is the probability of getting a king on the first draw times the probability of getting a king on the second draw given that you already got a king on the first. In symbols:  $\Pr(K_1 \& K_2) = \Pr(K_1) \times \Pr(K_2 | K_1)$ .

Now that we have conditional probabilities, we can *define* independence quite precisely.  $A$  and  $B$  are independent just in case the truth (or occurrence) of one has no influence or effect on the occurrence of the other; this means that

**Independence:**  $A$  and  $B$  are *independent* just in case  $\Pr(A) = \Pr(A | B)$ .

Whether  $B$  occurs (or is true) or not has no effect on whether  $A$  occurs (or is true). If we learn that  $B$  is true (or false), that should do nothing to change our beliefs about the probability of  $A$ .

Rule 5 tells us that *if*  $A$  and  $B$  are independent, then  $\Pr(A \& B) = \Pr(A) \times \Pr(B)$ . This is just a special case of the more general Rule 8. It works because *if*  $A$  and  $B$  are independent,  $\Pr(B) = \Pr(B | A)$ . So instead of writing  $\Pr(B | A)$  in the special case (independent conjuncts) covered by Rule 5, we can get by with the simpler  $\Pr(B)$ .

Rule 8. tells us that  $\Pr(A \& B) = \Pr(A) \times \Pr(B | A)$ . But we know that the order of the conjuncts in a conjunction doesn't affect the meaning of the conjunction:  $A \& B$  says the same thing as  $B \& A$ . So  $\Pr(A \& B) = \Pr(B \& A)$ . This means that  $\Pr(A \& B) = \Pr(B \& A) = \Pr(B) \times \Pr(A | B)$ . The value for this will be the same as the value we get when we use Rule 8, though in some cases one approach will be easier to calculate and in other cases the other one will be.

## 14.2 Analyzing Probability Problems

You must know the rules if you are to calculate probabilities.

### Summary of the Rules for Calculating Probabilities

1. **Events that are Certain to Occur:** If  $A$  is certain to be true,  $\Pr(A) = 1$ .
2. **Events that are Certain *not* to Occur:** If  $A$  is certain to be false,  $\Pr(A) = 0$ .
3. **Negations:**  $\Pr(\neg A) = 1 - \Pr(A)$ .
4. **Disjunctions with Incompatible Disjuncts:** If  $A$  and  $B$  are incompatible,  
 $\Pr(A \text{ or } B) = \Pr(A) + \Pr(B)$ .
5. **Conjunctions with Independent Conjuncts:** If  $A$  and  $B$  are independent,  
 $\Pr(A \& B) = \Pr(A) \times \Pr(B)$ .

**6. Disjunctions:**  $\Pr(A \text{ or } B) = \Pr(A) + \Pr(B) - \Pr(A \& B)$ .

**7. Definition of Conditional Probability:**  $\Pr(A|B) = \Pr(A \& B)/\Pr(B)$ .

**8. Conjunctions:**  $\Pr(A \& B) = \Pr(A) \times \Pr(B|A)$ .

### How to Approach a Problem

The key is to analyze a problem *before* you begin writing things down. The first question to ask yourself is: Am I calculating the probability of a negation, a disjunction, or a conjunction? The answer to this will tell you which rule is relevant to the problem; if you get this right, you are well on your way to a successful solution. In a complicated problem, you may have to use several of these rules in your calculations, but always begin by asking which rule applies first.

Begin by asking yourself the following questions.

1. If the sentence is a **negation**, use Rule 3..
  - Find the probability of the sentence that is being denied and subtract it from 1.
2. If it is a **disjunction**
  - Are the disjuncts incompatible? If so, use Rule 4.
  - Are the disjuncts compatible? If so, use Rule 6.
3. It is a **conjunction**?
  - Are the conjuncts independent? If so, use Rule 5.
  - Are the conjuncts dependent (= not independent)? If so, use Rule 8.

The tree diagram in Figure 14.3 on the next page (p. 256) represents the same information pictorially.

#### 14.2.1 Examples of Problem Analysis

*Problem A.* Suppose that you have a standard deck of 52 cards. You will draw a single card from the deck. What is the probability of drawing either an ace or a jack? Analysis of the problem:

1. You want to know about the probability of drawing an ace *or* drawing a jack, so you have a *disjunction*. The first disjunct is “I get an ace,” and the second disjunct is “I get a jack.” We could symbolize this as  $(A \text{ or } J)$ .
2. Are the disjuncts incompatible? Well, if you draw an ace you cannot also draw a jack (on that same draw). Getting an ace *excludes* getting a jack (and getting a jack excludes getting an ace). So the disjuncts are incompatible, and you use R4 (the rule of disjunctions with incompatible disjuncts).




Figure 14.3: Tree Diagram of Probability Problem Analysis

3. The rule says to *add* the probabilities of the two disjuncts:  $\Pr(A \text{ or } J) = \Pr(A) + \Pr(J)$
4. There are exactly four aces out of 52 cards, so  $\Pr(A)$  (the probability of drawing an ace) is  $4/52$  (which reduces to  $1/13$ ). There are also four jacks, so  $\Pr(J)$  is the same as that of drawing an ace, namely  $1/13$ .
5. Rule 4 tells us to add these probabilities:  $\Pr(A \text{ or } J) = 1/13 + 1/13 (= 2/13)$ .

*Problem B.* Suppose that you have a standard deck of 52 cards. You will draw a single card from the deck. What is the probability of drawing either a jack or a heart? Analysis:

1. You want to know about drawing a jack *or* drawing a heart, so you again have a *disjunction*. The first disjunct is “I get a jack” and the second disjunct is “I get a heart.” We symbolize this as  $(J \text{ or } H)$ .
2. Since you have a disjunction, the relevant rule will be one of the two Disjunction Rules. Which one it is depends on whether the disjuncts are incompatible.
3. Are the disjuncts incompatible? Well, if you draw a jack does that exclude drawing a heart? No. You might draw the jack of hearts. So the disjuncts are *not* incompatible, and you must use Rule 6 (the general disjunction rule).

4. This rule says to add the probabilities to the two disjuncts, but then “subtract out the overlap.” In other words, you must subtract out the probability that you get both a jack and a heart, and this is just the probability of getting the jack of hearts. So we have  $\Pr(J \text{ or } H) = \Pr(J) + \Pr(H) - \Pr(J \& H)$ .
5. There are four jacks out of 52 cards, so  $\Pr(J)$ , the probability of drawing a jack, is  $4/52$ . And there are 13 hearts, so  $\Pr(H)$ , the probability of drawing a heart, is  $13/52$ . Finally, there is just one possibility for getting a jack and a heart, namely the jack of hearts, so  $\Pr(J \& H)$  is  $1/52$ .
6. The General Disjunction Rule then tells us  $\Pr(J \text{ or } H) = 4/52 + 13/52 - 1/52$  (we won't worry about actually calculating such things until we get the basic concepts down—and even then you can use a calculator).

### Exercises

1. The chances of there being two bombs on a plane are very small, So when I fly I always take along a bomb. —Laurie Anderson.

What should we make of Anderson's advice (in light of the things we have learned thus far)?

2. What is the numerical value of  $\Pr(A|A)$ ? Explain why your answer is correct.
3. Suppose you are going to flip a fair coin. Which of the possible sequences is (or are) the most likely?
  1. HHHHTTTT
  2. HTHTHTHT
  3. HTHHTHTH
  4. HTHHTHTHT
  5. no one of these is any more likely than the others.
4. Suppose that you are about to turn over four cards from the top of a standard deck. Which of the following series of cards (in the order given) is the most likely?
  1. ace of hearts, king of diamonds, queen of spades, jack of hearts
  2. ace of heart, king of hearts, queen of hearts, jack of hearts
  3. ace of hearts, eight of spades, jack of diamonds, four of clubs
  4. no one of these is any more likely than the others.
5. If two sentences are incompatible, then
  1. they must also be independent.

2. the truth of one is completely irrelevant to the truth of the other.
  3. they cannot also be independent.
  4. none of the above.
6. You have an ordinary deck of 52 cards. You will draw a card, lay it on the table, then draw another card. (it is important to use the rules in these calculations)?
1. What is the probability of two kings?
  2. What is the probability of a queen on the second draw given a king on the first?
  3. What is the probability of a king on the first draw?
  4. What is the probability of the king of spades and the king of hearts (in either order)?
  5. What is the probability of a king and a queen?
  6. What is the probability of the jack of Diamonds and a spade (where the order in which you get the two doesn't matter)?
  7. What is the probability of not drawing a five at all?
7. Suppose that I am planning what to do this coming weekend, and the weather forecast is for 40% chance of rain on Saturday and 40% chance of rain on Sunday (40% chance = .40 probability). What is the probability that it will rain *sometime* or other during the weekend (assume that its raining or not on Saturday won't make it any more or less likely to rain on Sunday)?
8. What is  $\Pr(S|S)$ ? What about  $\Pr(S|\neg S)$ . Explain and defend your answers.
9. You and your friend Wilbur are taking a multiple-choice exam (and you are working independently and your answers are independent). There is exactly one correct answer to each question, and your task is to select it from five possible answers, 'a', 'b', 'c', 'd', and 'e'. You get to the third question and have no idea what the correct answer is, and the same thing happens with Wilbur. You guess 'a', and Wilbur guesses 'c'.
1. What is the probability that at least one of you got the correct answer?
  2. What is the probability that neither of you got the correct answer [the answer here is *not*  $4/5 \times 4/5$ ].
  3. You also guessed on the fourth problem. What is the probability that you got at least one of your two guesses is right?
10. Most automobile accidents occur close to home. Why do you suppose this is true? How could you explain what is involved using the notion of conditional probabilities?



## 14.3 Odds and Ends

In many situations we translate probabilities into the **odds** for or against a given outcome. For example, the probability of rolling a two when you roll a fair die is  $1/6$ , and the probability of not getting a two is  $5/6$ . We say that the odds of rolling a two are 1 to 5 and the odds against it are 5 to 1. The odds in favor of a two is the ratio of the numbers of ways of getting a two (one way) to the number of ways of not getting a two (five ways). And the odds against throwing a two are the five chances that some other side will come up against the one chance that a two will come up.

The relationship between odds and probabilities is a simple and straightforward:  $\Pr(A) = m/n$  if and only if the odds in favor of  $A$  are  $m$  to  $n - m$ .

### From Probabilities to Odds

If the probability of something is  $1/36$  (as is the probability of rolling box cars), then the odds in favor of it are 1 to 35 and the odds against it are 35 to 1. We convert probabilities to odds with the following rule: if the probability of a given outcome is  $m/n$ , then the odds in favor of it are  $m$  to  $n - m$  and the odds against it are  $n - m$  to  $m$ .

### From Odds to Probabilities

If your friend says that the odds of OU's beating Texas A&M are 1 : 5, what does she think the probability of A&M's winning is? We get the denominator for this probability by adding the two numbers in this ratio, so the number on the bottom is 6. Your friend believes that there is one chance in 6 that OU will win, which translates into a probability of  $1/6$ . And she also believes that the probability of A&M's winning is  $5/6$ . If the odds in favor of  $S$  are  $m$  to  $n$ , then the probability of  $S$  is the first number ( $m$ ) over the sum of the first and second numbers ( $m + n$ ).

### Fair Bets

Fair bets are based on the odds. If you want to make a *fair bet* that a two will come up when you roll a fair die, you should bet \$1 that you will get a two and your opponent should bet \$5 that you won't. If you both always bet these amounts, then over the long run you will both tend to break even. Gamblers call such a bet an *even-up proposition*.

By contrast, if you were to bet \$1 that you would roll a two and your opponent bets \$6 that you won't, then over the long haul you will come out ahead. And if

you bet \$1 that you will roll a 2 and your opponent bets \$4 that you won't, then over the long haul you will lose.

Organized gambling usually involves bets that are not even-up. A casino could not pay its operating expenses, much less turn a profit, if it made even-up bets. The house takes a percentage, which means paying winners less than the actual odds would require. The same is true for insurance premiums. It is also true for state lotteries, which in fact offer far worse odds than most casinos. If you gamble in such settings long enough, you are virtually certain to lose more than you win. Of course if you enjoy gambling enough you may be willing to accept reasonable losses as the price of getting to gamble.

### Example: Roulette

Roulette is a gambling game in which a wheel is spun in one direction and a ball is thrown around the rim in to the wheel in the opposite direction. A roulette wheel has a number of compartments, and players bet on which compartment the ball will land in. In the U.S. roulette wheels have thirty-eight compartments. They are numbered from 1 through 36; there is also a thirty-seventh compartment numbered 0 and a thirty-eighth numbered 00.

There are various bets players can place, but here we will focus on the simplest one where she bets that the ball will land on one specific number (say 14) from 1 through 36. Although the game can be complex, the following discussion gives the basic points.



Since there are thirty-eight compartments on the wheel, the probability that the ball will land on any given number, say 14, is  $1/38$ ;  $\Pr(14) = 1/38$ . Hence the *true odds* against rolling a 14 are 37 to 1. If you played the game over and over, betting at these odds, you would break even. You would win once every thirty-eight times, and the casino (the “house” or “the bank”) would win the other thirty-seven times. But when you did win, they would pay you \$37, which would exactly compensate

you for the thirty-seven times that you lost \$1 ( $37 \times \$1 = \$37$ ). We say that your bet has an expected value \$0.0.

But of course the house does not pay off at the true odds of 37 to 1. Instead

the house odds or betting odds against rolling a 14 are 35 to 1 (the house has the advantage of the 0 and 00). When you lose this doesn't make any difference. But when you win you get only \$36 (the \$35 plus the original \$1 that you bet). This is \$2 less than you would get if you were payed off at the true odds of 1 to 37. Since the house keeps \$2 out of every \$38 that would be payed out at the true odds, their percentage in  $2/38$ , or 5.26%. All but one of the bets you can make at roulette costs you 5.26% over the long haul (the remaining bet is even worse, from the player's point of view).

If you play just a few times, you may well win. Indeed, a few people will win over a reasonably long run. But the basic fact is that your bet on 14 has a negative *expected payoff* of -5.26 percent. This means that over the long run you will almost certainly lose at roulette. The odds are against you, and there are no systems or strategies or tricks that can change this basic fact. Simply put, there is absolutely no way you can expect to win at this game. There are a few *highly* skilled people who make a living playing poker, blackjack, or betting on the horses. But no one can make a living playing casino games like keno, craps, or roulette.

### Exercises on Odds and Probabilities

Calculate the odds and probabilities in each of the following cases.

1. What are the odds against drawing a king of spades from a full deck of playing cards?
2. What are the odds against drawing a king from a full deck?
3. What are the odds against drawing a face card from a full deck?
4. What are the odds against drawing a king if you have already drawn two cards (one a king, the other a six)?
5. You have a bent coin. The odds of flipping a head are 3 to 2. What is the probability of tossing a tail?
6. In Europe a roulette wheel only has thirty-seven compartments, one through 36 plus 0. They pay off at the same odds as U.S. casinos. How would this change the probabilities and the odds?
7. If the probability of Duke winning the NCAA basketball tournament championship is  $0.166 (= 1/6)$ , what are the odds that they will win? What are the odds against their winning? What are the fair bets for and against their winning? Defend your answers.

### 14.3.1 Sample Problems with Answers

In each case explain which rules are relevant to the problem. Your analysis of the problem is more important than the exact number you come up with.

1. You are going to roll a single die. What is the probability of rolling a two or an odd number?

1. You are asked about the probability of a *disjunction*; What is  $\Pr(T \text{ or } O)$ ?
2. Are the two disjuncts incompatible?
3. Yes. So we can use the simple disjunction rule (R4).
4. It says that  $\Pr(T \text{ or } O) = \Pr(T) + \Pr(O)$ .
5. And  $\Pr(T) + \Pr(O) = 1/6 + 3/6 = 4/6 (= 2/3)$ .

2. You are going to draw a single card from a full deck. What is the probability of getting either a spade or a three?

1. You are asked about the probability of a *disjunction*; What is  $\Pr(S \text{ or } T)$ ?
2. Are the two disjuncts incompatible?
3. No. They overlap because of the three of spades.
4. So we must use the more complex disjunction rule (R6) in which we subtract out the overlap.
5. It says that  $\Pr(S \text{ or } T) = \Pr(S) + \Pr(T) - \Pr(T \& S)$
6.  $\Pr(T \& S)$  is just the probability of drawing the three of spades, which is  $1/52$ .
7. So  $\Pr(S \text{ or } T) = \Pr(S) + \Pr(T) - \Pr(T \& S) = (13/52 + 4/52) - 1/52$

3. You are going to draw a two cards from a full deck *without* replacing the first card. What is the probability of getting exactly one king and exactly one queen (the order doesn't matter)?

1. You are asked about  $\Pr(K \& Q)$ , where order doesn't matter.
2. There are two different ways for this to occur:
  - (a) King on first draw and queen on second:  $(K_1 \& Q_2)$
  - (b) Queen on first draw and king on second:  $(Q_1 \& K_2)$
3. So we have to calculate the probability of a *disjunction*:  
What is  $\Pr[(K_1 \& Q_2) \text{ or } (Q_1 \& K_2)]$ ?
4. The two disjuncts are incompatible, so we use the simple disjunction rule (R4)
5. But each disjunct is itself a conjunction, and the conjuncts of each conjunction are *not* independent.

6. First disjunct is:  $\Pr(K_1 \& Q_2)$ . The general rule (R8) for conjunctions tells us that  $\Pr(K_1 \& Q_2) = \Pr(K_1) \times \Pr(Q_2 | K_1)$ , which is  $4/52 \times 4/51$
7. Second disjunct is:  $\Pr(Q_1 \& K_2)$ . It works the same way:  $\Pr(Q_1 \& K_2) = \Pr(Q_1) \times \Pr(K_2 | Q_1)$ , which is also  $4/52 \times 4/51$
8. Now add the probabilities for each disjunct:  $(4/52 \times 4/51) + (4/52 \times 4/51)$ .

### 14.3.2 More Complex Problems

In the next module we will look at a number of real-life applications of probability. We conclude this module with several problems that are more complex than the ones we've dealt with thus far.

Probability theory was formalized in the 1650s. The Chevalier de Méré was a wealthy Parisian gambler. He had devised a dice game that was making him money. He would bet even money (betting odds of 1:1) that he could roll at least one six in four throws of a die. Eventually people got wise to this game and quit playing it, so he devised a new game in which he bet even money that he could roll at least one double six (a six on each die) in twenty-four rolls of a pair of dice. But over time he lost money with this bet.

Finally he asked his friend, the philosopher and mathematician Blaise Pascal (1623-1662), why this was so. Pascal (and Fermat, with whom he corresponded) worked out the theory of probability and used it to explain why the first game was profitable while the second one was not. Let's see how to solve the first problem (the second is left as an exercise).

What is the probability of rolling at least one six in four throws of a die?

Rolling at least one six means rolling a six on the first roll, or the second, or the third, or the fourth. But it is difficult to work the problem in this way because one must subtract out all of the relevant overlaps.

It is easiest to approach this by way of its negation. The negation of the statement that you roll at least one six is the statement that you do not roll any sixes. This negation is equivalent to a conjunction: you do not roll a six on the first throw *and* you do not roll a six on the second throw *and* you do not roll a six on the third throw *and* you do not roll a six on the fourth throw. This conjunction has four conjuncts, but that doesn't really change anything that affects the probabilities. Each conjunct says that you get something other than a six, and so each has a probability of  $5/6$ .

Furthermore, each conjunct is independent of the other three (the die doesn't remember earlier outcomes). So we just multiply the probabilities of the four conjuncts to get the probability that the conjunction itself is true: the probability that

you don't get a six on any of the four rolls is  $5/6 \times 5/6 \times 5/6 \times 5/6 (= (5/6)^4)$ , which turns out to be  $625/1296$ .

This is the probability that you don't get any sixes. So the probability we originally asked about (getting at least one six) is just one minus this: the probability of getting at least one six is  $1 - (625/1296)$  (which is approximately  $671/1296$ ). This is just a bit more than  $1/2$ . This means that the odds of getting at least one six are 671 to 625, so over the long run the house will come out ahead, and you will lose if you keep playing their game.

## 14.4 Chapter Exercises

Exercises are included in most of the sections. Here we present some more difficult problems, extras for experts, although you now know enough to work at least some of the problems here. Answers to some of them are given below, but think about the problems before looking (you will need a pretty good calculator to get the exact numbers; if you don't have one, just work out the formulas).<sup>2</sup>

1. The probability that you will get a car for graduation is  $1/3$  and the probability that you will get a new computer is  $1/5$ , but you certainly won't get both. What is the probability that you will get one or the other?
2. You have a 35% chance of getting an *A* in Critical Reasoning and a 40% chance of getting an *A* in sociology. Does it matter whether the two outcomes are independent when you want to calculate the probability of at least one *A*? Does it matter whether the two outcomes are independent when you want to calculate the probability getting an *A* in both courses? Are they likely to be independent? Why?
3. Five of the 20 apples in the crate are rotten. If you pull out two at random, not replacing them as you pull them out, what's the probability that both will be rotten?

The remaining problems are harder.

4. Now solve the Chevalier de Méré's second problem. What is the probability of rolling at least one double six in twenty four rolls of a pair of dice? Use the same strategy that was used above to solve his first problem.
5. *Aces and Kings* Remove all of the cards except the aces and kings from a deck. This leaves you with an eight-card deck: four aces and four kings. From this deck, deal two cards to a friend.



<sup>2</sup>Some of these are classic problems. Problem 3. is due to Steve Selvin in 1975. Answer to some of the problems, along with further references, will be found in Chapter 27.

1. If he looks at his cards and tells you (truthfully) that his hand contains an ace, what is the probability that both of his cards are aces?
2. If he instead tells you (truthfully) that one of his cards is the ace of spades, what is the probability that both of his cards are aces?

The probabilities in the two cases are *not* the same.

6. *It's in the Bag* There are two non-transparent bags in front of you. One contains two twenty dollar bills and the other contains one twenty and one five. You reach into one of the bags and pull out a twenty. What is the probability that the other bill in that bag is also a twenty?

7. *The Monty Hall Problem* There are three doors in front of you. There is nothing worth having behind two of them, but there is a suitcase containing \$50,000 behind the third. If you pick the right door, the money is yours.



You choose door number 1. But before Monty Hall shows you what is behind that door, he opens one of the other two doors, picking one he knows has nothing behind it. Suppose he opens door number 2. This takes 2 out of the running, so the only question now is about door 1 and door 3.



You may now reconsider your earlier choice: you can either stick with door 1 or switch to door 3.

1. What is the probability that the money is behind door 1?
2. What is the probability that the money is behind door 3?
3. Do your chances of winning improve if you switch?

8. *The Birthday Problem* How many people would need to be in a room for there to be a probability of .5 that two of them have a common birthday (born on the same day of the month, but not necessarily of the same year)? Assume that a person is just as likely to be born on any one day as another and ignore leap years.



Hint: Much as in the previous problem, it is easiest to use the rule for negations in answering this.

### Answers to Selected Problems

7. *The Monty Hall Problem* We will work the answer out in a later chapter using the rules for calculating probabilities. For now, here are three hints (don't look at the third until you have tried working the problem). First, you do improve your chances by switching to door 3? Second, think about how what would happen if you repeated this process a hundred times. Third, draw a diagram representing all of the things that could happen and not how often switching pays off compared to the total number of outcomes.

8. *The Birthday Problem.* The negation of the claim that at least two people in the room share a birthday is the claim that none of them share a birthday. If we can calculate the latter, we can subtract it from 1 to get the former.

Order the people by age. The youngest person was born on one of the 365 days of the year. Now go to the next person. She could have been born on any of the 365 days of the year, so the probability that her birthday *differs* from that of the first person is  $364/365$ . Now move on to the next person. The probability that his birthday differs from those of the first and the second is  $363/365$ . For the next person, the relevant probability is  $362/365$ , and so on.

The birthdays are independent of one another, so the probability that the first four people have different birthdays is  $365/365 \times 364/365 \times 363/365 \times 362/365$ . There is a pattern here that we can generalize. The probability that the first  $n$  people have different birthdays is  $365 \times 364 \times \cdots \times (365 - (n - 1))/365^n$ . And so the probability that at least two out of  $N$  people have a common birthday is one minus all of this, i.e.,  $(1 - (365 \times 364 \times \cdots \times (365 - (N - 1)))/365^N)$

Now that we have this formula, we can see what values it gives for different numbers of people (and so for different values of  $N$ ). When there are twenty two people in the room  $N$  is 22, and the formula tells us that the probability that at least two of them have a common birthday is about .47. For twenty three people it is slightly more than a half (.507). For thirty two people, the probability of a common birthday is over .75, and for fifty people it is .97. And with one hundred people there is only about one chance in three million that none share a common birthday.



**14.4.1 Summary of Rules for Calculating Probabilities**

- 1. Events that are Certain to Occur:** If  $A$  is certain to be true,  $\Pr(A) = 1$ .
- 2. Events that are Certain *not* to Occur:** If  $A$  is certain to be false,  $\Pr(A) = 0$ .
- 3. Negations:**  $\Pr(\neg A) = 1 - \Pr(A)$ .
- 4. Disjunctions with Incompatible Disjuncts:** If  $A$  and  $B$  are incompatible,  
 $\Pr(A \text{ or } B) = \Pr(A) + \Pr(B)$ .
- 5. Conjunctions with Independent Conjuncts:** If  $A$  and  $B$  are independent,  
 $\Pr(A \ \& \ B) = \Pr(A) \times \Pr(B)$ .
- 6. Disjunctions:**  $\Pr(A \text{ or } B) = \Pr(A) + \Pr(B) - \Pr(A \ \& \ B)$ .
- 7. Definition of Conditional Probability:**  $\Pr(A | B) = \Pr(A \ \& \ B) / \Pr(B)$ .
- 8. Conjunctions:**  $\Pr(A \ \& \ B) = \Pr(A) \times \Pr(B | A)$ .



## **Part VI**

# **Induction in the Real World**



## **Part IV. Induction in the Real World**

In this part we will examine several ways that induction works in the real world.

In Chapter 15 we learn about a few notions from descriptive statistics; you need to understand them to interpret many of the things you will read outside of class. We will then look at samples and populations and some of the ways in which we draw conclusions about populations from premises about samples. We will conclude with a look at correlations.

In Chapter 16 we turn to various applications of probabilistic notions. We will examine the notion of expected value and several other applications of probability. We will then examine several ways in which our probabilistic reasoning often goes wrong in daily life; here we will examine the gambler's fallacy, the conjunction fallacy, regression to the mean, and some common mistakes about coincidence.



# Chapter 15

## Samples and Correlations

**Overview:** We begin this chapter with a few basic notions from descriptive statistics; you need to understand them to interpret many of the things you will read outside of class. We will then look at samples and populations and some of the ways in which we draw conclusions about populations from premises based on samples. We then consider correlations. These have to do with the degree to which various things are related. Correlations underwrite many of our predictions, but we are often mistaken about the degree to which things are correlated.

### Contents

---

<b>15.1 Descriptive Statistics</b> . . . . .	<b>274</b>
15.1.1 Features of Samples . . . . .	275
15.1.2 Exercises . . . . .	277
<b>15.2 Inferences from Samples to Populations</b> . . . . .	<b>278</b>
15.2.1 Sampling in Every day Life . . . . .	278
15.2.2 Samples and Inference . . . . .	279
15.2.3 Good Samples . . . . .	279
15.2.4 Bad Sampling and Bad Reasoning . . . . .	283
15.2.5 Exercises . . . . .	284
<b>15.3 Correlation</b> . . . . .	<b>287</b>
15.3.1 Correlation is Comparative . . . . .	289
15.3.2 Exercises . . . . .	294
<b>15.4 Real vs. Illusory Correlations</b> . . . . .	<b>295</b>
15.4.1 Ferreting out Illusory Correlations . . . . .	297

15.4.2 The Halo Effect: A Case Study in Illusory Correlation	299
<b>15.5 Chapter Exercises</b> . . . . .	<b>301</b>

In New York City in the summer a number of cats fall from open windows in high-rise apartment buildings. On August 22, 1989 *The New York Times* reported the startling fact that cats that fell further seemed to have a better chance of survival. When they checked with the Animal Medical Center, the paper found that 129 cats that had fallen were brought in for treatment. Seventeen of these were put to sleep by their owners (in most cases because they could not afford treatment, rather than because the cat was likely to die). Eight of the remaining 115 cats died. But the surprising thing is that the cats that fell the furthest seemed to have the highest probability of living. Only one of the 22 cats that fell from above 7 stories died, and there was but a single fracture among the 13 that fell more than 9 stories. What could account for this?

## 15.1 Descriptive Statistics

We will begin with several basic concepts from **descriptive statistics** that are important for reasoning. We won't be concerned with formulas for calculating them, but you will encounter these concepts outside of this class, so you need to learn what they mean.

**Sample:** subgroup of a population

A **population** is a group of things (e.g., Oklahoma voters, households, married couples, fruit flies). And a **sample** is a subgroup of the population. For example, we might conduct a poll of 1000 college graduates and ask them how much they earned. These 1000 people would constitute our sample; the parent population would be *all* college graduates. In the next section we will see that information about samples can be used to draw inferences about entire populations, but in this section we will be concerned with description rather than inference.

**Parameter:** a feature of a population

A **parameter** is some numerical characteristic of an entire population (e.g., average GPA of all OU students, average income of all college graduates; as we will see in a moment, it could also be a measure of dispersion or a measure of correlation. For example, an average income in the population of adult U. S. citizens of \$18,525 is a parameter.

**Statistic:** a feature of a sample

By contrast, a **statistic** is a corresponding numerical characteristic of a sample (e.g., the average GPA of the OU students contacted in a recent survey). One



way to remember what goes with what is that the two *p*-words—*p*opulation and *p*arameter—go together, and the two *s*-words—*s*ample and *s*tatistic—go together.

### 15.1.1 Features of Samples

Properties or characteristics that come in degrees are called **variables**. For example, the age, weight, and income of people in the United States are variables. Each of them can take on many different values: Wilbur weighs 165 pounds, Martha 103, and Sam 321  $\frac{1}{4}$ . We can also think of more abstract things as variables; for example, probability is a variable that can take any of the infinitely many values from 0 to 1. In the simplest case a variable might only have two values; for example, sex is a variable with the values male and female (such variables are important; they are called *dichotomous variables*).

**Variables:** properties or features that come in degrees

When the members of a population or sample are measured with respect to some variable like their score on the ACT test, the resulting set of all the numerical scores is a **distribution** of values for that variable. Thus, the set of all the ACTs scores last year is a distribution of values for the variable of 1998 ACT scores. Similarly, the set of all the scores on the first exam in this class is a distribution of the variable of scores on the first exam.

It can be difficult to see what a large distribution of values really amounts to; we get lost in a sea of numbers. So it is often useful to *condense* the information in the distribution into simpler numbers. The most basic ways of doing this is to calculate measures of central tendency. There are three common measures of this sort.

#### Measures of central tendency

The **mean** is what you already know under the name *average*. To find the mean of a distribution you add all of the numbers in the distribution together and divide by the number of items in the distribution. When the class gets an exam back, the first thing many people want to know is the average (i.e., mean) score on the test; this tells them how well the class did as a whole. The mean is the most important measure of central tendency, but it has the weakness that it is affected by just a few extreme values.

**Mean:** the average

The **median** of a distribution is the number such that half the numbers in the distribution are less than it and half are greater. The median of the numbers 1, 2, 3, 4, 5 is 3, because two numbers are less than it and two are greater. What if no single number splits a distribution into two equal parts, as occurs in the distribution 1, 2, 3, 4? Here we will take the number half way between 2 and 3, i.e., 2.5 as the median; clearly half the cases fall below it and half fall above.

**Median:** splits the group into halves

**Mode:** most frequent value(s)

The **mode** of a distribution is the value that occurs most frequently in it. The mode of 1, 2, 3, 2, 4 is 2, because 2 occurs twice and all the other numbers occur only once. A distribution may have more than one mode. For example, the distribution 1, 2, 3, 2, 4, 4, 2, 4 has two modes: 2 and 4.

What are the mean, median, and mode of the following set of numbers: 179, 193, 99, 311, 194, 194, 179?

1. **Mean:** Add the seven numbers together, which yields 1349. Then divide this by 7, which (rounding off) comes to 192.7
2. **Median:** The median is easiest to see if we list these numbers in order of magnitude, as 99, 179, 179, 193, 194, 194, 311. Here we find that 193 splits the distribution into two equal parts, so it is the median.
3. **Mode:** This distribution has two numbers which occur twice, 179 and 194. So it has two modes, 179 and 194.

### Measures of Dispersion

Measures of central tendency are often useful. For example, it will help you understand how you did on an exam to know the class average (the mean). And it will be easier to choose a major if you know the average number of people with that major who found jobs soon after they graduated. But measures of central tendency don't tell us much about the relative position of any given item or about the extent to which values are spread out around a mean.

For example, the distributions

- 7, 8, 8, 9
- 1, 3, 11, 17

have the same mean, namely 8. But the items in the first distribution are clustered much more tightly around the mean than those of the second. If the values in a distribution are quite spread out, then the mean may not be very informative.

**Measures of dispersion** provide additional information; they tell us how spread out ("dispersed") the values in a distribution are.

The **range** is the distance between the largest and the smallest value in the distribution. In the distribution: 179, 193, 99, 311, 193, 194, 179, the range is the distance between 311 and 99, i.e.,  $311 - 99 = 212$

**Percentiles** Often a numerical value or score doesn't tell you much in and of itself. If you learn that you scored a 685 on the math component of the ACT or that you got an 86 on the first exam in this course, that doesn't really tell you how well you did. What you want to know is how well you did in comparison with those who took the same exam. Percentiles provide information about such relative positions.

The **percentile rank** of a value or score is the percentage of values that fall below it. For example, if Sandra got an 86 on the first exam and 75% of the class got lower grades, than Sandra's score has a percentile rank of 75%. And her score, 86, falls at, or *is*, the *75th percentile*.

Percentiles provide relative positions in percentage terms. For example, suppose that 100 people take the first exam and that Wilbur gets a 79. If 60 (= 60%) students scored lower than 79, then Wilbur's score of 79 falls at the 60th percentile.

**Quartiles** work like the median. The first quartile is the value such that 1/4 of the values are less than it, the second quartile the value such that half of the values are less than it (this number is also the median), the third quartile the value such that 3/4 of the values are less than it. The first quartile falls at the 25th percentile.

The **standard deviation** is a very important measure of dispersal. We can't actually calculate it without a formula (which we won't worry about here), but the intuitive idea is that the standard deviation measures the *average distance of all the values from the mean*. It tells us how far, on average, the values deviate from the mean or average value in the distribution. The greater the standard deviation, the more spread out the values are. Hence, although the distributions 7, 8, 8, 9 and 1, 3, 11, 17 have the same mean, namely 8, the first will have a lower standard deviation than the second.

### 15.1.2 Exercises

1. Find the mean, median, mode, and range of each of the following distributions (which we may think of as measurements of people's weight in pounds):

1. 176, 132, 221, 187, 132, 194, 190
2. 176, 193, 99.5, 321, 112, 200, 120

Here is a list of people in a class, their score on their final, and the percentage of people who scored below them. In each case, give the percentile where their grade falls.

1. Sheila got a 97, 95% scored lower
2. Bruce got a 46, 5% scored lower
3. Wilbur got a 85, 80% scored lower

3. Which distribution will have the greater standard deviation:

- 10, 11, 14, 9
- 6, 9.5, 10, 18.67

## 15.2 Inferences from Samples to Populations

We frequently use sample statistics to draw *inductive inferences* about population parameters. When a newspaper conducts a poll to see how many people think Clinton should be impeached, they check with a sample, say 2000 adults across the U. S. and draw a *conclusion* about what American adults in general think. Their results would be more accurate if they checked with everyone, but when a population is large, it simply isn't practical to examine all of its members. We have no choice but to rely on a *sample* from the population and make an inference based on it. When scientists engage in such inferences they are said to be using *inferential statistics*. But all of us draw inferences from samples to populations many times every day.

### 15.2.1 Sampling in Every day Life

Inferences based on samples are common in medical research, the social sciences, and polling. In these settings scientists use what are called *inferential statistics* to move *from* claims about samples *to* conclusions about populations.

But we all draw similar inferences many times each day. You are driving through Belleville for the first time and trying to decide where to eat. You have had good experiences at McDonald's restaurants in the past (the set of McDonald's restaurants where you have eaten in the past at constitutes your sample). So you might conclude that all McDonald's restaurants (the population) are likely to be good and decide to sample the culinary delights of the one in Belleville.

Or suppose you know six people (this is your sample) who have dated Wilbur, and all of them found him boring. You may well conclude that almost everyone (this is the population) would find him boring.

Whenever you make a generalization on the basis of several (but not all) of the cases, you involved in sampling. You are drawing a conclusion about some larger group on the basis of what you've observed about one of its subgroups.

### Learning

Most *learning from experience* involves drawing inferences about unobserved cases (populations) from information about a limited number of cases that we have observed (our samples). You know what strategies worked in the cases you have experienced (your sample) for getting a date, quitting smoking, getting your car running when the battery seems dead, or doing well on an exam. And you then draw conclusions about the relevant populations on the basis of this knowledge.

An examination is really a sampling procedure to determine how much you have learned. When your calculus professor makes up an examination, she hopes to sample items from the population of things you have learned in her course and to use your grade as an indicator of how much information you have acquired.

### 15.2.2 Samples and Inference

We often infer a conclusion about a population from a description of a sample that was drawn from it. When we do

1. Our **premises** are claims about the **sample**.
2. Our **conclusion** is a claim about the **population**.




Figure 15.1: Inference from Sample to Population

For example, we might draw a conclusion about the divorce rate of people living in Oklahoma from premises describing the divorce rates among 800 couples that the OU Human Relations Department sampled.

In such a case our inference is *not* deductively valid. It involves an inductive leap. The conclusion goes beyond the information in the argument's premises, because it contains information about the *entire population* while the premises only contain information about the *sample*. But if we are careful, our inference can still be inductively strong. This means that if we begin with true premises (which in this case means a correct description of the sample), we are likely to arrive at a true conclusion (about the entire population).

### 15.2.3 Good Samples

A good inductive inference from a sample to a population requires:

1. A **large** enough sample

## 2. A **representative** (unbiased) sample

### Good samples:

1. Big enough
2. Representative

We would need to delve more deeply into probability to say exactly how large is large enough, but we won't need to worry about that here. The important point is that in everyday life we very often rely on samples that are clearly too small.

We also need a sample that is as *representative of the entire population* as possible. A sample that is not representative is said to be *biased*. An unbiased sample is *typical* of the population. By contrast, in a biased sample, some portions of the population are *overrepresented* and others are *underrepresented*.

The problem with a very small sample is that it is not likely to be representative. Other things being equal, a bigger sample will be more representative. But there are costs to gathering information—a price in time, dollars, and energy—so it is rarely feasible to get huge samples.

We can never be certain that a sample is unbiased, but we can strive to avoid any non-trivial biases we can discover. With some thought, the worst biases are often fairly obvious. Suppose, for example, that we want to know what the adult public as a whole (our population) thinks about the consumption of alcohol. We would clearly get a biased sample if we distributed questionnaires only at a pool hall (we would have an overrepresentation of drinkers and an underrepresentation of those favoring temperance) or only at the local meetings of MADD (here the biases would be reversed).

A classic example of a biased sample occurred in 1936 when a magazine, *The Literary Digest*, conducted a poll using names in telephone directories and on car registration lists. A majority of the people they sampled favored Alf Landon over Franklin Roosevelt in that year's Presidential election, but when election day rolled around Roosevelt won in a landslide. What went wrong? News organizations now use telephone polls routinely, but in 1936 a relatively small percentage of people had telephones and cars, and most of them were affluent. These were the people most likely to vote for the Republican candidate, Landon, and so the sample was not representative of all voters.

There are other cases where bias is likely, even though it won't be this blatant. For example, anytime members of the sample volunteer, e.g., by returning a questionnaire in the mail, we are likely to have a biased sample. People willing to return a questionnaire are likely to differ in various ways from people who are not. Or, to take an example closer to home, tests that focus on only some of the material covered in class are likely to elicit a biased, unrepresentative sample of what you have learned. They aren't a fair sample of what you know. Unfortunately, in some cases biases may be difficult to detect, and it may require a good deal of expertise to find it at all.

### Random Sampling

The best way to obtain an unbiased sample is to use random sampling. A **random sample** is a method of sampling in which each member of the population has an equally good chance of being chosen for the sample. Random sampling does not guarantee a representative sample—nothing short of checking the entire population can guarantee that—but it does make it more likely. Random sampling avoids the biases involved with many other methods of sampling.

We can rarely get a truly random sample, but under some conditions, typically in carefully conducted studies and surveys, we can come reasonably close. But even in daily life we can use samples that are much less biased than those we often rely on.

**Random sample:** each member of population has equal chance of being sampled

### Random Digit Dialing (RDD)

Modern technology now allows national polling organizations with large resources to approach the ideal of random sampling. Polls like the *New York Times*/CBS News poll use what is called *Random Digit Dialing* (RDD). The goal here is to give every residential phone number an equal chance of being called for an interview. Nowadays almost all major polls use some form of RDD. For example, the *New York Times*/CBS News poll uses the GENESYS system, which employs a database of over 42,000 residential telephone numbers throughout the U.S. that is updated every few months. The system also employs software that draws a random sample of phone numbers from this database and then randomly makes up the last four digits of the number to be called. Of course some sorts of people are harder to reach on the phone than others, and some sorts are more willing to volunteer information over the phone. But RDD constitutes an impressive step in the direction of randomization.<sup>1</sup>

Scientists sometimes go a step further and use a *stratified random sample*. Here the aim is to ensure that there is a certain percentage of members of various subpopulations in our sample (e.g., an equal number of men and of women). We separate the population into relevant categories or “strata” before sampling (e.g., into the categories or subpopulations of men and of women). Then we sample randomly within each category.

With the GENESYS system, the breakdown of people contacted in a nationwide poll is:

---

<sup>1</sup>So pollsters have learned a great deal since the *Literary Digest* poll in 1936, and their polls are now much more accurate. There are other problems with surveys and polls; for example, the responses we get often depend in subtle ways on the way that our questions are worded. We will study these matters in a later chapter when we discuss framing effects. A discussion of the GENESYS system and of the breakdown of groups noted below may be found in the *New York Times*, 11/4/99.

1. 22 % from the northeast, 33 % from the south, 24 % from the mid-west, and 21 % from the west.
2. 47 % men and 53 % women.
3. 80 % white, 11 % black, 1 % Asian, and 6 % other.
4. 24 % college graduates, 27 % with some college or trade schooling, 37 % who did not go beyond high school, and 12 % who did not graduate from high school.
5. As of 1999, 27 % who consider themselves Republicans, 36 % who consider themselves as Democrats, and 30 % who consider themselves Independents.
6. 20 % who consider themselves liberals, 42 % who consider themselves moderates, and 32 % who consider themselves conservatives.

### Polls

A growing practical problem in recent years has been the decline in the public's participation in polls. Pollsters are getting more and more "nonresponses." Some of these result from the difficulty in contacting people by phone (people are at work, away somewhere else, on a dial-up connection to the internet, etc.). But those contacted are also less willing to participate than they were in the past. The reasons for this aren't completely clear, but growing disillusionment with politics and lack of patience for unsolicited calls resulting from the increase in telemarketing may be part of the reason.

The use of *push polling* also leads to a wariness about polls. Push polling is not really polling at all. Instead, an organization, e.g., the campaign organization of a candidate running for Senate, calls thousands of homes. The caller says he is conducting a poll, but in fact no results are collected and instead a damaging—and often false—claim about the other side is implanted in what is presented as a neutral question. For example, the caller might ask: "Do you agree with Candidate X's goal to cut social security payments over the next six years?" Such deceptive uses of polling is likely to make the public more cynical about polls.

### Sampling Variability

It is unlikely that any sample that is substantially smaller than the parent population will be *perfectly* representative of the population. Suppose that we drew many samples of the same size from the same population. For example, suppose that we drew many samples of 1000 from the entire population of Oklahoma voters. The set of all of these samples is called a *sampling distribution*.

The samples in our sampling distribution will vary from one to another. This just means that if we draw many samples from the same population, we are likely to



get somewhat different results each time. For example, if we examine twenty samples, each with 1000 Oklahoma voters, we are likely to find different percentages of Republicans in each of our samples.

This variation among samples is called **sampling variability** or **sampling error** (though it is not really an error or mistake). Suppose that we just take one sample of 1000 Oklahoma voters and discover that 60% of them prefer the Republican candidate for Governor. Because of sampling variability, we know that if we had drawn a different sample of the same size, we would probably have gotten a somewhat different percentage of people favoring the Republican. So we won't be able to conclude that exactly 60% of all Oklahoma voters favor the Republican from the fact that 60% of the voters in our single sample do.

As our samples become larger, it actually becomes *less* likely that the sample mean will be *exactly* the same as the mean of the parent population. But it becomes *more* likely that the mean of the sample will be *close* to the mean of the population. Thus, if 60% of the voters in a sample of 10 favor the Republican candidate, we can't be very confident in predicting that about 60% of voters in general do. If 60% of a sample of 100 do, we can be more confident, and if 60% of a sample of 1000 do, we can be more confident still.

Statisticians overcome the problem of sampling variability by calculating a **margin of error**. This is a number that tells us how close the result of a poll should usually be to the population parameter in question. For example, our claim might be that 60% of the population, *plus or minus three percent*, will vote Republican this coming year. The smaller the sample, the larger the margin of error. But there are often large costs in obtaining a large sample, so we must compromise between what's feasible and margin of error. Here you do get what you pay for.

It is surprising, but the size of the sample does not need to be a large percentage of the population in order for a poll or survey to be a good one. What is important is that the sample not be biased and that it be large enough; once this is achieved, it is the absolute number of things in the sample (rather than the proportion of the population that the sample makes up) that is relevant for taking a reliable poll.

In our daily life we can't hope for random samples, but with a little care we can avoid flagrant biases in our samples, and this can improve our reasoning dramatically.

### 15.2.4 Bad Sampling and Bad Reasoning

Many of the reasoning errors we will study in later chapters result from the use of small or biased samples. Drawing conclusions about a general population on the basis of a sample from it is often called *generalization*. And drawing such a conclusion from a sample that is too small is sometimes called the *fallacy of hasty*

*generalization* or, in everyday language, *jumping to a conclusion*.

Our samples are also often biased. For example, a person in a job interview may do unusually well (they are striving to create a good first impression) or unusually poorly (they may be very nervous). If so, his actions constitute a biased sample, and they will not be an accurate predictor of his performance on the job. In coming chapters we will find many examples of bad reasoning that result from use of samples that are too small, highly biased, or both.

### **Example: The Two Hospitals**

There are two hospitals in Smudsville. About 50 babies are born everyday in the larger one, and about 14 are born everyday in the smaller one down the street. On average 50% of the births in both hospitals are girls and 50% are boys, but the number bounces around some from one day to the next in both hospitals.

- Why would the percentage of boys vary from day to day?
- Which hospital, if either, is more likely to have more days per year when over 65% of the babies born are boys?

We all know that bigger samples are likely to be more representative of their parent populations. But we often fail to realize that we are dealing with a problem that involves this principle; we don't "code" it as a problem involving sample size.

Since about half of all births are boys and about half are girls, the true percentages in the general population are about half and half. Since a smaller sample will be less likely to reflect these true proportions, the smaller hospital is more likely to have more days per year when over 65% of the births are boys. The births at the smaller hospital constitute a smaller sample.

This will seem more intuitive if you think about the following example. If you flipped a fair coin four times, you wouldn't be all that surprised if you got four heads. The sample (four flips) is small, so this wouldn't be too surprising; the probability is  $1/16$ . But if you flipped the same coin one hundred times, you would be very surprised to get all heads. A sample this large is very unlikely to deviate so much from the population of all possible flips of the coin.

### **15.2.5 Exercises**

#### **Exercises**

1. What is the probability of getting all heads when you flip a fair coin (a) four times, (b) ten times, (c) twenty times, (d) one hundred times?

In 2–7 say (a) what the relevant population is, (b) what the sample is (i.e., what is in the sample), (c) whether the sample seems to be biased; then (d) evaluate the inference.

2. We can't afford to carefully check all of the computer chips that come off the assembly line in our factory. But we check one out of every 300, using a randomizing device to pick the ones we look at. If we find a couple of bad ones, we check out the whole bunch.

3. We can't afford to carefully check all of the computer chips coming off the assembly line. But we check the first few from the beginning of each day's run. If we find a couple of bad ones, we check out the whole bunch.

4. There are more than a hundred nuclear power plants operating in the United States and Western Europe today. Each of them has operated for a number of years without killing anybody. So nuclear power plants don't seem to pose much of a danger to human life.

5. Joining the Weight Away program is a good way to lose weight. My friend Millie and my uncle Wilbur both lost all the weight they wanted to lose, and they've kept it off for six months now.

6. Joining the Weight Away program is a good way to lose weight. *Consumer Reports* did a pretty exhaustive study, surveying hundreds of people, and found that it worked better than any alternative method for losing weight. [This is a fictitious example; I don't actually know the success rates of various weight-reduction programs.]

7. Alfred has a good track record as a source of information about the social lives of people you know, so you conclude that he's a reliable source of gossip in general.

8. Pollsters are getting more and more "nonreponses." Some of these result from the difficulty in contacting people by phone (people are at work, away somewhere else, on a dial-up connection to the internet, etc.). But those contacted are also less willing to participate than they were in the past. Does this automatically mean that recent polls are more likely to be biased? How could we determine whether there were. In which ways might they be biased (defend your answer).

9. We began the chapter with the story of cats that had fallen from windows (p. 274). What could account for these finding?

10. Several groups of studies have shown that the survival rates of patients after surgery is actually lower in better hospitals than in ones that aren't as good. More people die in the better hospitals. What might account for this?

11. The psychologist Robyn Dawes recounts a story about his involvement with a committee that was trying to set guidelines for professional psychologists. They were trying to decide what the rules should be for reporting a client (and thus breaching confidentiality) who admitted having sexually abused children when the abuse occurred in the distant past. A number of people on the committee said that they should be required to report it, even when it occurred in the past, because the one sure thing about child abusers is that they never stop on their own without professional help. Dawes asked the others how they knew this. They replied, quite sincerely, that as counselors they had extensive contact with child abusers. Does this give us good reason to think that child abusers rarely stop on their own? What problems, if any, are involved in the group's reasoning?

12. A poll with a 4% margin of error finds that 47% of voters surveyed plan to vote for Smith for Sheriff and 51% plan to vote for her opponent Jones. What things would you like to know in evaluating these results? Assuming that sound polling procedures were used, what can we conclude about who will win the election?

13. Suppose that a polling organization was worried that a telephone poll would result in a biased sample. So instead they pick addresses randomly and visit the residences in person. They have limited resources, however, and so they can only make one visit to each residence, and if no one is home they just go on to the next place on their list. How representative is this sample likely to be? What are its flaws? Could they be corrected? At what cost?

14. Suppose that you are an excellent chess player and that Wilbur is good, but not as good as you. Would you be more likely to beat him in a best of three series or in a best of seven series (or would the number of games make any difference)? Defend your answer.

15. A recent study conducted by Centers for Disease Control and Prevention was reported in the *Journal of the American Medical Association* (JAMA) (CNN, 11/3/99). The study was based on a survey of 9,215 adult patients who were members of Kaiser Permanente health maintenance organization (an HMO). The study focused on eight childhood traumas, including psychological, physical and sexual abuse, having a battered mother, parents who are separated or divorced, and living with those were substance abusers, mentally ill or had been imprisoned.

It was found that more traumas seemed to increase the probability of smoking. For example, people who experienced five or more traumas were 5.4 times more likely to start smoking by the age of 14 than were people who reported no childhood trauma. Analyze this study. What was the sample and the population. What are possible strong points and weak points?

In 16 and 17 explain what (if anything) is wrong in each of the following examples. When reasoning goes wrong, it often goes very wrong, so it's quite possible for even a short argument to be flawed in more than one way.

16. Suppose that the NRA (National Rifle Association) recently conducted a poll of its members and found that they were overwhelmingly opposed to any further gun control measures.

17. Suppose that a random sample of 500 OU students showed that an overwhelming majority do not think that OU needs to budget more money for making the campus more accessible to handicapped people.

18. A noted TV psychologist recently noted that the average marriage lasts 7 years, a fact that she sought to explain by pointing to evidence that life goes in 7 year cycles (so that it was not surprising that marriages should last 7 years). How may she have misunderstood the fact that she was explaining?

#### Answers to Selected Exercises

10. The actual cause, though you could only hypothesize it on the basis of the information in this problem, turns out to be that high-risk patients, those needing the most dangerous types of surgery, often go to better hospitals (which are more likely to provide such surgery, or at least more likely to provide it at a lower risk).

14. Hint: think about the hospital example above.

### 15.3 Correlation

Some variables tend to be related. Taller people tend to weigh more than shorter people. People with more education tend to earn more than people with less. Smokers tend to have more heart attacks than non-smokers. There are exceptions, but “on average” these claims are true.

There are many cases where we want to know the *extent* to which two variables are related. What is the relationship between the number of cigarettes someone smokes and his chances of getting lung cancer? Is there some relationship between years of schooling and average adult income? What is the connection between class attendance and grades in this course? Learning the answers to such questions is important for discovering how to achieve our goals (“Since the chances of getting cancer go up a lot, I’ll try to quit smoking even though I really enjoy it.”).

**Correlation** is a measure of the degree to which two variables are related—the degree to which they vary together (“covary”). If two things tend to go together,

**Correlation:** the degree to which two variables are related

then there is a *positive correlation* between them. For example, the height and weight of people are positively correlated; in general, greater height means greater weight. On the other hand, if two things tend to vary inversely there is a *negative correlation* between them. For example, years of schooling and days spent in prison are negatively correlated; in general, more years of schooling means less time in jail. And if two things are completely unrelated, they are not correlated at all.

Correlations underwrite predictions

Correlations between variables are extremely important in *prediction*. If you knew the heights of all of the students in my critical reasoning class last semester, you would be able to make more accurate predications about each student's weight than if you didn't know their heights. You would still make some mistakes, but on average your predictions would be more accurate.

There is a formula for calculating correlations, and the resulting values are numbers between +1.0 (for complete positive correlation) and -1.0 (for a complete negative correlation); a correlation of 0 means that there is no pattern of relationship between the two variables. This allows for very precise talk about correlations. We won't worry about such precision here, however, but will simply focus on the basic ideas.

### Correlation and Probability

We could apply the things we learned about probability to cover all cases of correlation, but here we will just get the general idea by considering the case of two dichotomous variables (variables that only have two values).

Consider the smoking variable and its two values, smoker and non-smoker, and the heart-attack variable and its two values, having a heart attack and not having a heart attack. The two variables are not independent. Smokers are more likely than non-smokers to have heart attacks, so there is a positive correlation between smoking and heart attacks. This means that  $\Pr(H|S) > \Pr(H) > \Pr(H|\neg S)$ . Or in words, the property of having a heart occurs at a higher rate in one group (smokers) than in another group (people in general, as well as the group of people who don't smoke). So correlation compares the rate at which a property (like having a heart attack) occurs in two different groups.

Zero correlation = Independence

If the correlation were negative, we would instead have  $\Pr(H|S) < \Pr(H)$ . And if there were no correlation at all, the two variables would be independent of each other, i.e.,  $\Pr(H|S) = \Pr(H)$ . Correlation is *symmetrical*. That means that it is a two-way street. If  $S$  is positively correlated with  $H$ , then  $H$  is positively correlated with  $S$ , and similarly for negative correlations and for non-correlations. In terms of probabilities this means that if  $\Pr(A|B) > \Pr(A)$ , then  $\Pr(B|A) > \Pr(B)$  (exercise for experts: prove this).

### 15.3.1 Correlation is Comparative

The claim that there is a positive correlation between smoking and having a heart attack does *not* mean that a smoker is highly likely to have a heart attack. It does *not even* mean that a smoker is more likely than not to have a heart attack. Most people won't have heart attacks even if they do smoke. The claim that there is a positive correlation between smoking and having a heart attack instead simply means that there are more heart attack victims among smokers than among non-smokers.

		<b>Heart Attack</b>	
		+	-
<b>Smoke</b>	+	12	80
	-	7	90

Figure 15.2: Thinking about Correlations

A good way to get a rough idea about the correlation between two variables is to fill in some numbers in the table in Figure 15.2. It has four cells. The + means the presence of a feature (smoking, having a heart attack) and the - means not having it (being a non-smoker, not having a heart attack). So the cell at the upper left represents people who are both smokers and suffer heart attacks, the cell at the lower left represents people who are non-smokers but get heart attacks anyway, and so on. We could then do a survey and fill in numbers in each of the four cells.

The key point to remember is that smoking and heart attacks are correlated just in case  $\Pr(S|H) > \Pr(S|\neg H)$ . So you cannot determine whether or not they are correlated merely by looking at  $\Pr(S|H)$ . This number might be high simply

Correlation is comparative




Figure 15.3: Correlation between Smoking and Heart Attacks

because the probability of suffering a heart attack is high for everyone, smokers and nonsmokers alike. Correlation is *comparative*: you have to compare  $\Pr(S|H)$  to  $\Pr(S|\neg H)$  to determine whether smoking and heart attacks are correlated or not.

### Comparative Diagrams to Illustrate Correlation

One of the easiest way to understand the basics of correlation is to use a diagram like that in Figure 15.3. Diagrams like this are more rough and ready than the diagram above, but they are easier to draw. The percentages here are hypothetical and are simply used for purposes of illustration. Here we suppose that the percentage of smokers who suffer heart attacks is 30% and that the percentage of nonsmokers who suffer heart attacks is 20% (these round numbers are chosen to make the example easier; they are not the actual percentages).

In this *comparative diagram* the horizontal line in the smokers column indicates that 30% of all smokers suffer heart attacks and the lower horizontal line in the nonsmokers column indicates that 20% of nonsmokers suffer heart attacks.

The fact that the percentage line is *higher* in the smokers column than it is in the nonsmokers column indicates that smoking is positive correlation between being a smoker and having a heart attack. It is the *relationship* between these two horizontal lines that signifies a positive correlation. Similarly, the fact that the percentage line is *lower* in the nonsmokers column indicates that there is a negative correlation between being a nonsmoker and having a heart attack.

The further apart the lines are in a diagram like this, the stronger the correlation is. So Figure 15.4 on the facing page illustrates an even stronger positive correlation between smoking and heart attacks.






Figure 15.4: A Stronger Positive Correlation

Finally, if the lines were instead the same height, say at 30% (as in Figure 15.5), smoking and having a heart attack would be independent of one another: they would not be correlated, either positively or negatively.




Figure 15.5: Independence between Smoking and Heart Attacks

Notice that to draw such diagrams you do not need to know exact percentages. You only need to know which column should have the higher percentage, i.e., the higher horizontal line.

### Correlation and Causation

Correlations often point to causes; they are *evidence* for claims about what causes what. When two variables, like smoking and having a heart attack, covary we suspect that there must be some reason for their correlation—surely *something* must cause them to go together. But correlation is not the same thing as causation. For one thing, correlation is symmetrical (smoking and heart attacks are correlated with each other), but causation is a one-way street (smoking causes heart attacks, but heart attacks rarely cause people to smoke). So just finding a positive correlation doesn't tell us what causes what.

Correlation  $\neq$  Causation

When your child's pediatrician says "Spots like this usually mean measles," she is relying on a positive correlation between the presence of spots and having the measles. We know the spots don't cause the measles, and commonsense suggests that measles causes the spots. But sometimes variables are correlated with each other even when neither has any causal influence on the other. For example, every spring my eyes start to itch and a day or two later I have bouts of sneezing. But the itchy eyes don't cause the sneezing; these two things are joint effect of a third thing, allergies to pollen, that causes both of them.

There are many examples of correlations between things that are effects of some third, common cause. The scores of identical twins reared in very different environments are correlated on a number of behavioral variables like introversion–extroversion. If the twins were separated at birth and reared apart, one twin's high degree of extroversion cannot be the cause of the other's extroversion. In this case their high degrees of extroversion are joint effects of a third thing—a **common cause**—namely having the same genotype (genetic makeup).




Figure 15.6: Common Causes

There are many examples of correlations that are based on a common cause. Every fall my eyes become itchy and I start to sneeze. There is a high positive correlation between these two things, but neither causes the other. Both are the joint effects of a common cause: pollen to which I'm allergic (Figure 15.6). Similarly, there is a positive correlation between a falling barometer and a rain storm, but

neither causes the other. They are both caused by an approaching cold front. So sometimes variables are correlated because they have a common cause rather than because either causes the other.

Some early spokesmen (they were all men in those days) for the tobacco companies tried to convince the public that something similar was true in the case of smoking. They urged that smoking and heart attacks are correlated because they are *common effects* of some third thing. Some peoples' genetic makeup, the spokesmen suggested, led them to smoke *and also* made them more susceptible to heart disease. Despite much research, a common genetic cause for smoking and cancer was never found, but the research was necessary to exclude this possibility. We can never rule out the possibility of common causes without empirical observations.

In many cases it is difficult to determine what causes what, even when we know a lot about correlations. For example, in the late 1990s, the rate of violent crime in many U. S. cities dropped. The drop is accompanied by a number of things, e.g., more police on the beat, tougher sentencing laws, various educational programs. Thus there is a (negative) correlation between number of police and number of crimes, between tougher sentences and number of crimes (more police, less crime), and so on.

But there is a great deal of debate about just what *caused* the drop in crime (naturally everyone involved wants to take credit for it). Of course it may be that each of these things, e.g., more police, plays some causal role. It is very difficult to determine just how much difference each of the factors makes, but we need to do so if we are going to implement effective measures to reduce crime.

It is also known that self-esteem and depression are negatively correlated. Lower self-esteem tends to go with depression. But what causes what? Lower self-esteem might well lead to depression, but depression might also lower self-esteem. Of there could be a vicious circle here, where each condition worsens the other. But it is also possible that there is some third cause, e.g., a low level of neurotransmitters in the brain or negative events in one's life.

As these examples show, finding causes is often important for solving serious problems like crime or depression. But while correlations can frequently be detected by careful observation, tracking down causes is often much more difficult. It is best done in an experimental setting, where we can control for the influence of the relevant variables.

### Correlation and Inferential Statistics

Once we determine whether two variables are correlated in a sample we may want to draw inferences about whether they are correlated in the population. Here the material earlier in this chapter on inferential statistics is relevant.

### 15.3.2 Exercises

1. Say whether the correlation between the following pairs of variables is strong, moderate, weak, and in those cases that do not involve dichotomous variables, say whether the correlation is positive or negative. Defend your answer (if you aren't sure about the answer, explain what additional information you would need to discover it); in each case think of the numbers as measuring features of adults in the US:

1. height and weight
2. weight and height
3. weight and caloric intake
4. weight and income
5. weight and score on the ACT
6. weight and amount of exercise,
7. weight and gender
8. years of schooling and income

2. Having schizophrenia and being from a dysfunctional family are positively correlated. List several possible causes for this correlation. What tests might determine which possible causes are really operating here?

3. How might you determine whether watching television shows depicting violence and committing violent acts are correlated in children under ten? Suppose that they were: what possible causes might explain this correlation?

4. Many criminals come from broken homes (homes where the parents are separated or divorced). Explain in detail what you would need to know to determine whether there really is a correlation between being a criminal and coming from a broken home. Then explain what more you would need to know to have any sound opinion on whether coming from a broken home causes people to become criminals.

5. How would you go about assessing the claim that there is a fairly strong positive correlation between smoking marijuana and getting in trouble with the law?

6. We often hear about the power of positive thinking, and how people who have a good, positive attitude have a better chance of recovering from many serious illnesses. What claim does this make about correlations? How would you go about assessing this claim?

7. Many criminals come from broken homes (homes where the parents are separated or divorced). Explain in detail what you would need to know to determine

whether there really is a correlation between being a criminal and coming from a broken home. Then explain what more you would need to know to have any sound opinion on whether coming from a broken home causes people to become criminals.

8. Suppose that 70% conclude from this about the evils of marijuana? Explain your answer.

9. Suppose that 30% of those who smoke marijuana get in trouble with the law, and 70% do not. Suppose further that 27% of those who don't smoke marijuana get in trouble with the law and 73% do not. What are the values of  $\Pr(T|M)$  and  $\Pr(T|\neg M)$ . Are smoking marijuana and getting in trouble with the law correlated? If so, is the correlation positive or negative? Does it seem to be large or small?

10. Suppose we obtain the following statistics for Wilbur's high school graduation class: 46 of the students (this is the actual number of students, not a percentage) who smoked marijuana get in trouble with the law, and 98 did not. And 112 of those who didn't smoke marijuana got in trouble with the law and 199 did not. What are the values of  $\Pr(T|M)$  and  $\Pr(T|\neg M)$ ? Are smoking marijuana and getting in trouble with the law correlated? If so, is the correlation positive or negative? Does it seem to be large or small?

11. Suppose that last year the highway patrol in a nearby state reported the following: 10 people who died in automobile accidents were wearing their seatbelts and 37 were not wearing them. Furthermore, 209 people who did not die (but were involved) in accidents were wearing their seatbelts, while 143 were not wearing them. Does this give some evidence that seatbelts prevent death in the case of an accident? Is there a non-zero correlation between wearing seat belts and being killed in an accident? If so, is it positive or negative, and what is the relative size (large, moderate, small)? Be sure to justify your answers.

Extras for Experts. Prove that positive correlation is symmetrical. That is, prove that  $\Pr(A|B) > \Pr(A)$  just in case  $\Pr(B|A) > \Pr(B)$ .

## 15.4 Real vs. Illusory Correlations

A pitfall that is especially relevant to this chapter is belief in illusory correlations. We believe in an **illusory correlation** when we think we perceive a correlation where one doesn't really exist. More generally, we believe in an illusory correlation when we think that things go together *substantially* more (or less) often than they actually do.

**Illusory correlation:** something that looks like a correlation, but isn't

A recurrent theme in this course is that human beings are constantly seeking to explain the world around them. We look for order and patterns, and we tend

to “see” them even when they don’t exist. For example, most of us will think we detect patterns in the random outcomes of flips of a fair coin. So it is not surprising that we tend to see strong relationships—correlations—among variables even when the actual correlation between them is minimal or nonexistent. This can be a serious error, because once we think we have found a correlation we typically use it to make predictions and we frequently develop a causal explanation for it. If the correlation is illusory, the predictions will be unwarranted and our explanation of it will be false.

If Wilbur, for example, believes that women tend to be bad drivers—i.e., if he thinks there is a correlation between sex and driving ability—then it will be natural for him to predict that he will encounter more bad drivers among women than among men. He may even go so far as to predict that Sue, whose driving he has never observed, will be a bad driver. Finally, he may look around for some explanation of why women don’t drive well, one that may suggest they don’t do other things well either. So beliefs in illusory correlations have consequences, and they are typically bad.

Our tendency to believe in illusory correlations has been verified repeatedly in the lab. In a series of studies in the 1960s, Loren and Jean Chapman gave subjects information that was supposedly about a group of mental patients. The subjects were given a clinical diagnosis of each patient and a drawing of a figure attributed to the patient. The diagnoses and drawings, which were all fictitious, were constructed so that there would be no correlation between salient pairs of features; for example, the figure was just as likely to have weird eyes when the diagnosis was paranoia as when it wasn’t.

Subjects were then asked to judge how frequently a particular diagnosis, e.g., paranoia, went along with a particular feature of the drawing, e.g., weird eyes. Subjects greatly overestimated the extent to which such things went together, i.e., they overestimated the correlation between them, even when there was data that contradicted their conclusions. And they also had trouble detecting correlations that really *were* present.

Various things lead us to think we detect correlations when none exist. As we would by now expect, context and expectations often play a major role. We have some tendency to see what we expect, and even hope, to see. And we have a similar tendency to find the patterns we expect, and even hope, to find. For example, in word association experiments, subjects were presented with pairs of words (‘tiger - bacon’, ‘lion - tiger’).

They later judged that words like ‘tiger’ and ‘lion’, or ‘bacon’ and ‘eggs’, which they would expect to go together, had been paired much more frequently than they actually had been. Similarly, if Wilbur expects to encounter women who are bad drivers, he is more likely to notice those who do drive badly, forget about

those who don't, and interpret the behavior of some good women drivers as bad driving.

Many beliefs in illusory correlation amount to superstitions. If I believe that my psychic friends down at the psychic hotline accurately predict the future, then I believe that there is a positive correlation between what they say and what turns out to be true (i.e., I believe that the probability that a prediction will be true, given that they say it will, is high). Again, we may remember cases where someone wore their lucky sweater and did well on the big exam, so we see an (illusory) correlation between wearing it and success.

Illusory correlations often arise in our reasoning about other people. Many of us tend to think that certain good qualities (like honesty and kindness) are correlated, so when we learn that a person has one good feature, we think it more likely that she has others. In some cases she may, but it's not reasonable to draw this conclusion without further evidence. This pattern of thinking occurs so frequently that it has a name—the *halo effect*—and we return to it in more detail near the end of this chapter.

Illusory correlations also make it easier for people to cling to stereotypes. A *stereotype* is an oversimplified generalization about the traits or behavior of the members of some group. It attributes the same features to *all* members of the group, whatever their differences. There are many reasons why people hold stereotypes, but belief in illusory correlations often reinforces them. Thus people may believe that members of some race or ethnic group tend to have some characteristic—usually some negative characteristic like being lazy or dishonest—which is just to say that they believe that there is a correlation between race and personality traits.

But even when our expectations and biases don't color our thinking, we often judge that two things go together more often than they really do simply because we *ignore* evidence to the contrary. It is often easier to think of positive cases where two things go together than to think of negative cases where they don't.

Suppose we learn about several people who have the same illness and some of them got better after they started taking Vitamin E. It can be very tempting to conclude that people who take Vitamin E are more apt to recover than those who do not. But this may be an illusory correlation. Perhaps they would have gotten better anyway—people often do. To know whether there is a genuine correlation here, we need to compare the recovery rate among those who took the vitamin and those who did not.

Illusory correlations often result from overemphasis on positive cases

### 15.4.1 Ferreting out Illusory Correlations

In later chapters we will learn to guard against many of the things that encourage belief in illusory correlations, but we are already in a position to note one very

important remedy. In this example we were inclined to see a correlation between taking Vitamin E and recovering from the illness because we focused on just one sort of case, that where people took the vitamin *and* got better. But many people who don't take the vitamin may also recover, and perhaps many other people who do take it don't recover. In fact, it might even turn out that a higher percentage of people who don't take the vitamin get better. *Correlation is comparative.*

One way to begin to see the importance of other cases is to note that the case of people who don't take Vitamin E but recover anyway provides a *baseline* against which we can assess the effectiveness of the vitamin. If 87% of those who don't take the vitamin recover quickly, then the fact that 87% of those who *do* take it recover quickly doesn't constitute a positive correlation between taking the vitamin and recovery. If 87% of those who *don't* take it recover quickly and if 86% (which sounds like a pretty impressive percentage, if we neglect the contrast cases) of those who do recover, taking the vitamin actually *lowers* the chances of recovery.

A more realistic example illustrates the same point. We may easily remember students who smoked marijuana and got into non-drug-related trouble with the law. They may stand out in our mind for various reasons, perhaps because they are frequently cited as bad examples. This can lead to belief in an illusory correlation between smoking dope and getting into trouble. It may well be that such a correlation exists, but to determine whether it does, we also have to consider the *contrast groups*. In other words, we have to consider not just group 1, but also groups 2, 3, and 4:

**Group 1** People who smoked marijuana and did get in trouble

**Group 2** People who smoked marijuana but did not get in trouble

**Group 3** People who didn't smoked marijuana and did get into trouble

**Group 4** People who did not smoked marijuana and did not get into trouble

The relevant question here is whether the probability of getting in trouble is higher if you smoke marijuana than if you don't. In other words, is it true that

$$\Pr(T|M) > \Pr(T|\neg M)?$$

And it is impossible to answer this question without considering all four groups.

To estimate a person's probability of getting in trouble *given* that she smoked marijuana ( $\Pr(T|M)$ ), we must first estimate the proportion of marijuana users who did get in trouble, which requires some idea about users who got in trouble (Group 1) and users who did not (Group 2). And then to estimate the probability of a person's getting in trouble given that they did not smoke marijuana ( $\Pr(T|\neg M)$ ), we need to estimate proportion of non-users who got in trouble, which requires



some idea about non-users who got in trouble (Group 3) and those who did not (Group 4).

But we tend to focus on cases where both variables, here smoking marijuana and getting in trouble with the law, are present. This is an example of our common tendency to look for evidence that confirms our hypotheses and or beliefs, and to overlook evidence that tells against them. This is called **confirmation bias**, and we will examine it in detail in a later chapter on testing and prediction. But for now the important point is that we can only make sensible judgments about correlations if we consider all four of the groups in the above list.

In real life we are unlikely to know exact percentages, and we won't usually bother to write out tables like the ones above. But if we have reasonable, ballpark estimates of the actual percentages, quickly constructing a comparative table in our heads will vastly improve our thinking about correlations. If we just pause to ask ourselves about the three cells we commonly overlook, we will avoid many illusory correlations. We will get some practice at this in the following exercises.

#### 15.4.2 The Halo Effect: A Case Study in Illusory Correlation

##### Seeing More Connections Than Are There

When we give a person a strong positive evaluation on one important trait (like intelligence), we often assume that they should also receive positive evaluations on other traits (like leadership potential). This is called the **halo effect**. The one positive trait sets up a positive aura or halo around the person that leads us to expect other positive traits.

**Halo effect:** if someone has one good trait we tend to jump to the conclusion that they have many others

The reverse also holds; when a person seems to have one important negative trait, we tend to think that he will have other negative traits as well. The halo effect is a common example of our vulnerability to illusory correlations. We tend to think that one trait (e.g., honesty) is highly correlated with another (e.g., courage), when in fact it may not be. We don't do this consciously, but it shows up in our actions.

In one real-world study, flight commanders tended to see a strong relationship between the intelligence of a flight cadet and his physique, between his intelligence and his leadership potential, and between his intelligence and his character. These traits are not completely unrelated, but the commanders greatly overestimated the strength of their connections. In another study, students who were told that their instructor would be warm were more likely to see him as considerate, good-natured, sociable, humorous, and humane. Being warm set up a halo that they thought extended to these other traits.

If two traits really do tend to go together, then we can draw a reasonable (but fallible) inference from one to the other. But such inferences are only legitimate

*if* there truly is a strong objective connection—a high correlation—between the two traits. In many cases there is not, so the halo effect leads us to “see” more correlations or connections than there really are. We tend to see sets of traits as package deals when in fact they are quite separate.

### **What is Beautiful is Good**

Physical attractiveness provides one of the most striking examples of the halo effect. Different cultures perceive different things as attractive, but within most cultures (or subcultures), there is a good deal of agreement on what is viewed as attractive and what is not. Many people act as though they believe that there is a strong positive correlation between physical attractiveness (as rated by members of their culture) and a large number of positive characteristics. For example, physically attractive people are seen as happier, stronger, kinder, and more sensitive than less attractive people.

Of course there may be some connection between being attractive and being happy or between being attractive and having good social skills (why might this be so?). But attractiveness creates a halo that extends to completely unrelated characteristics. For example, experimenters had subjects read a set of essays. Each essay had a picture attached to it that the experimenter said was a picture of the author (although this was just a ruse). The quality of an essay was judged to be better when it was attributed to an attractive author.

Illusory correlations based on attractiveness occur in many settings in the real world. Attractive job candidates are more likely to be hired than less attractive ones. In one real-world study, physically attractive men earned a higher starting salary, and they continued to earn more over a ten-year period, than less attractive men. And although physically attractive women did not have higher starting salaries, they soon earned more than their less attractive counterparts.

The phenomenon even affects basic issues involving justice and fairness. The transgressions of attractive children are judged less severely by adults than similar actions by less attractive children. A mock jury sentenced an unattractive defendant to more years in prison than an attractive defendant, even though the crime was described in exactly the same words in each case. And killing an attractive victim gained a stiffer sentence than killing an unattractive one.

Perhaps these findings should not be surprising. Beauty is held up as an ideal in commercials, movies, and TV, and on-screen heroes and heroines are almost always attractive. In fact, there is a *physical attractiveness stereotype*, and this is probably what sets up the halo. Once we classify someone as attractive, the attractiveness stereotype or schema is activated, and we find it natural to suppose that a person has other components of the stereotype.

There are a few exceptions to the attractiveness halo. Physically attractive women are more likely to be judged vain and egotistical, although people tend to think better of beautiful women, unless they are seen as misusing their beauty. Physically attractive men are more likely to be judged less intelligent. But in general physical attractiveness establishes a strong, positive halo.

As in most cases of the halo effect, the physical attractiveness stereotype is based on bad reasoning (although it does have some features of a *self-fulfilling prophecy*: if attractive people are *treated* better, they may *do* better in various ways). It is also unfair. But if we know about the phenomenon, we can more easily guard against it in our own judgments and try to protect ourselves against other people's tendencies to fall victim to it in their own reasoning.<sup>2</sup>

## 15.5 Chapter Exercises

1. I know people who think that there is a strong positive correlation between having the astrological sign Libra and being indecisive. Why might they have come to think this? What claim does this make about correlations? How would you go about assessing this claim?

2. The following passages are from a column in *The Oklahoma Daily* (April 13, 1998) in which Rep. Bill Graves argues that homosexuals and lesbians should not be employed as support personnel at public schools. For each of the following, (1) explain Graves' point in citing the statistics, and (2) critically evaluate his use of the statistics.

- "The 1948 Kinsey survey . . . found that 37 percent of homosexual men and 2 percent of lesbians admitted sexual relations with children under 17 years old. Twenty-eight percent of homosexual men and 1 percent of lesbians admitted sexual relations with children under 16 years old while they were age 18 or older."

---

<sup>2</sup>You can learn more about sampling and statistics from any good book on statistics, e.g., Larry E. Toothaker, *Introductory Statistics for the Behavioral Sciences*. McGraw-Hill, 1986. D. S. Moore's *Statistics: Concepts and Controversies*, NY: W. H. Freeman & Company, 1985 is another accessible book that includes a discussion of survey techniques. For an accessible account of the Chapmans' work see L. J. Chapman & J. Chapman, "Test Results are what you Think they are," *Psychology Today*, November 1971, pp. 18–22. The exercise about child abusers stopping on their own is from Robyn Dawes, *Rational Choice in an Uncertain World*, Harcourt Brace College Publishers, 1988, Ch. 6; this book contains excellent treatments of several of the topics discussed in this and subsequent chapters. The example of the falling cats is discussed by Marilyn Vos Savant in volume in *Prade Magazine*.

- "The average age of homosexual men is 39 years, and 45 years for lesbians. Thus, that lifestyle is actually a death style from which children should be protected."

3. Evaluate the following interchange:

*Wilma:* I've just graduate from law school and now I have to take the bar exam. I'm sort of nervous about passing.

*Wilbur:* Don't about 90% of those who take it pass?

*Wilma:* It's given twice a year. In the summer about 90% pass it; in the winter it's about 70%

*Wilbur:* I'd take it in the summer if I were you.

4. Political consultants increasingly use focus groups in an effort to determine which themes, even which words, their candidate should use to get more votes. What is a focus group and how do they work (heck the web if you aren't sure)? Then explain the ways in which the concepts introduced in this chapter, e.g., sample and population, bear on the use of such groups and the evaluation of their responses.

5. Personnel Director for a large company: We are very careful in our job interviews. We see some very good people and the decisions are often tough. But looking back, we have almost always made the best decisions. The people we have hired have worked out very well.

6. When a teacher gives you an examination, they are taking a sample of the things you have learned in the course. Explain, in more detail, what the sample and the population here are. What does it mean for a sample here to be biased? In what ways are good tests unbiased?

#### Answers to Selected Exercises

5. There is a problem with this sample. In order to see whether the hiring decisions were the best the Personnel Director would need to know how the people she *didn't* hire would have worked out. This is nearly impossible to know, but failing this it would be useful to know how the people who weren't hired ending up doing at the job they eventually got.

# Chapter 16

## Applications and Pitfalls

**Overview:** In this chapter we consider the notion of expected value and several other applications of probability. We then examine several ways in which our probabilistic reasoning often goes wrong in daily life; here we will examine the gambler's fallacy, the conjunction fallacy, regression to the mean, and some common mistakes about coincidence.

### Contents

---

<b>16.1 What do the Numbers Mean?</b> . . . . .	<b>304</b>
16.1.1 Ratios of Successes to Failures . . . . .	304
16.1.2 Frequencies . . . . .	305
16.1.3 Degrees of Belief . . . . .	305
16.1.4 How can we Comprehend such Tiny Numbers? . . . .	306
16.1.5 Probabilistic Reasoning without Numbers . . . . .	307
<b>16.2 Expected Value</b> . . . . .	<b>308</b>
16.2.1 Pascal's Wager . . . . .	310
<b>16.3 The Gambler's Fallacy</b> . . . . .	<b>311</b>
<b>16.4 The Conjunction Fallacy</b> . . . . .	<b>312</b>
<b>16.5 Doing Better by Using Frequencies</b> . . . . .	<b>314</b>
<b>16.6 Why Things go Wrong</b> . . . . .	<b>315</b>
<b>16.7 Regression to the Mean</b> . . . . .	<b>317</b>
16.7.1 Regression and Reasoning . . . . .	319
<b>16.8 Coincidence</b> . . . . .	<b>320</b>

---

 16.9 Chapter Exercises . . . . . 321
 

---

## 16.1 What do the Numbers Mean?

You can now calculate the probabilities of various things happening. But what do the numbers you get tell you—*what* do they *mean*? The answer is that they mean different things in different cases. We will note three important cases, and the answer is different for each of them.

### 16.1.1 Ratios of Successes to Failures

With common games of chance we can determine the probabilities of simpler outcomes intuitively; indeed, it is much easier to do this than it is to calculate them with the relevant rule. Let's analyze what we do when we make these intuitive determinations. You are going to draw one card from a full deck. What is the probability that you'll draw a king? You didn't need any complicated rules to answer this. Instead, you reason that: there are four kings out of fifty two cards, and we are equally likely to draw any one of them, so the probability of getting a king is  $4/52$ .

In such cases where each of the outcomes is equally likely to occur we take the number of outcomes of interest to us, divide it by the total of possible outcomes, and interpret this ratio as a probability.

$$\text{Number of outcomes of interest} = \frac{\text{number of outcomes of interest}}{\text{number of all possible outcomes}}$$

For example, in the case of drawing a king from a deck the outcomes of interest are getting a king, and there are four of these. And the set of all possible outcomes consists of drawing any of the fifty two cards in the deck. If we call the outcomes of interest a *success* (a terminology that goes back to the gambling roots of probability), we can say that the probability of a success is

$$\frac{\# \text{ Successes}}{\# \text{ Possible cases}}$$

A similar approach works for outcomes of drawing balls from an urn, throwing dice, spinning a roulette wheel, and the like.

But this approach only works when the basic cases of interest are *equally likely*. It works when we flip a *fair* coin; the probability of heads is the numbers of cases

of interest, or the number of cases of interest, *successes* as they are often called, over the number of possible cases. There is one way to flip a head and two possible outcomes. So the probability is  $1/2$ . But this doesn't work if we flip a biased coin, say one that is twice as likely to come up heads as tails. There is still just one way to have a success (i.e., to flip a heads) and just two possible outcomes (heads and tails), but the probability of a head will no longer be  $1/2$ . To handle cases like this, and many real-life cases as well, we need to turn to frequencies.

### 16.1.2 Frequencies

In many cases probabilities are empirically determined frequencies or proportions. For example, the probability that a teenage male driver will have an accident is the percentage or frequency of teenage male drivers who have accidents. This approach applies, though sometimes less clearly, to the price of insurance premiums, weather forecasting, medical diagnosis, medical treatment, divorce, and many other cases.

For example, a health insurance company records the frequency with which males over 50 have heart attacks. It then translates this into a probability that a male over 50 will have a heart attack, and charges accordingly. Again, when your doctor tells you that there is a 5% chance that a back operation will worsen your condition, she is basing her claim on the fact that about 5% of the people who get such operations get worse. The outcomes of interest (getting worse) divided by the total number of cases (all those having this kind of surgery) is  $5/100$ .<sup>1</sup>

In many cases some of the possible outcomes are more likely to occur than others, but we can adapt the basic approach by viewing the probability of a given sort of event as the relative frequency with which it occurs (or would occur) in the set of possible outcomes.

### 16.1.3 Degrees of Belief

Often we do not have access to solid information about frequencies, and in some cases it isn't even clear which frequencies are relevant. But even in these cases we often have beliefs that involve something very like probabilities. For example, I don't have solid information about frequencies that would let me assess the probability that aliens from outer space have infiltrated the OU golf team. Nevertheless, I believe that the probability that they have is very low. Or, to take a more serious example, if you serve on a jury you may have to form a judgment about the likelihood that the defendant is guilty.

---

<sup>1</sup>At least this is so if you got lucky in your choice of doctors; there is some evidence that many physicians' estimates of probabilities of cures are quite inaccurate.

It may be unclear how I can assign a probability to the statement: “Aliens from outer space have infiltrated the golf team” (let’s abbreviate this as  $A$ ). But whatever rough probability value I assign it, my beliefs will only cohere with each other if I assign further rough probabilities in accordance with the rules of probability. For example, since I think that the probability of  $A$  is very low, I believe that  $1 - \Pr(A)$  of its negation is very high. And I believe that  $\Pr(A \text{ or } \neg A) = 1$  and that  $\Pr(A \ \& \ \neg A) = 0$ .

In short, probabilities sometimes represent ratios involving equally likely cases, they sometimes represent frequencies, and they sometimes represent our degrees of belief. The former are much easier to work with, but many things that matter in life involve the second or third. Fortunately for us, these issues don’t matter a lot in the sorts of cases we are likely to encounter.<sup>2</sup>

#### 16.1.4 How can we Comprehend such Tiny Numbers?

We can develop some *feel* for the meaning of frequency probabilities when they aren’t too small. For example, the probability that you will roll a two with a fair die is  $1/6$ . This means that on average, over the long run, you will roll a two one-sixth of the time.

But many probabilities are much smaller numbers. For example, the probability of getting kings on two successive draws from a full deck when we replace the first card is  $4/52 \times 4/52$  (approximately .0059) whereas the probability of getting two kings when we don’t replace the first card is  $4/52 \times 3/51$  (approximately .0045). We aren’t used to thinking about such tiny numbers, and it is difficult to get a grip on what they mean. In a highly technological world, the differences between numbers like this are sometimes important, and they also matter to casinos that want to stay in business. But such differences don’t matter much to us in our daily life, and we won’t agonize over them. The important point for us is that most of us have a poor feel for very large and very small numbers even in cases where their relative sizes are *very* different.

We have considered the probabilities of outcomes when we draw cards or roll dice, but people also consider the probabilities of outcomes in cases that matter a lot more, including matters of life and death. What is the likelihood of dying in a plane crash? Of getting cancer if you smoke? Of contracting AIDS if you don’t use a condom?

Terrorism is frightening, and many people have canceled trips to Western Europe because of their perceptions of the terrorist threat. In fact, however, fewer than one in a million Americans are killed by terrorists in any given year, whereas

<sup>2</sup>There are other interpretations of the meaning of probability. The issues are controversial and can get pretty complex, but we don’t need to worry about such subtleties here.



over one in 5000 are killed in automobile accidents. The differences between the probabilities of these two occurrences is enormous, and any rational assessment of travel-related risks should take this into account.

If we had a good feel for large numbers, we could apply this to probability; for example, it would give us a better feel for the magnitude of the difference between  $1/5000$  and  $1/1,000,000$ . But most of us are no better with big numbers than with small ones. When we hear about the size of the national debt, which is measured in trillions of dollars, the numbers are so enormous that our minds just go numb.

A good way to develop some feel for the meanings of very large and very small numbers is to translate them into concrete terms, ideally into terms that we can visualize. What does one thousand really mean? What about ten thousand? Well, the Lloyd Nobel Center seats roughly ten thousand (11,100), and Owen Stadium seats roughly seventy thousand.

With larger numbers visualization becomes difficult, but analogies can still be useful. Consider the difference between one million (1,000,000) and one billion (1,000,000,000). It takes eleven and a half days for one million seconds to elapse, whereas it takes thirty two years for one billion seconds to tick away (how long does it take for one trillion—1,000,000,000,000—seconds to elapse?). And the relative difference in probabilities of one in a million and one in a billion is equally immense.

### Exercises

1. If it takes about 32 years for a billion seconds to elapse, how long does it take for a trillion seconds to elapse? Explain how you arrived at your answer.
2. How can we apply the points we have learned about the differences between a million and a billion to the claims that one alternative has a chance of one in a million of occurring and a second alternative has a chance of one in a billion of occurring?
3. OU's fieldhouse seats about 10,000 people and the football stadium seats about 75,000 people. Can you think of any concrete image that could help you get an intuitive handle on the number 1,000,000? Give it your best shot.

### 16.1.5 Probabilistic Reasoning without Numbers

In our daily lives we rarely worry about *precise* probability values; indeed, such numbers are often unattainable or even meaningless. But In the next few sections we will see that the concepts we acquired in mastering the rules of probability will help us understand many things that happen in real life. We will see how

probabilistic concepts are relevant even in the absence of precise numerical values for probabilities.

## 16.2 Expected Value

Most things in life are uncertain, so we don't have any choice but to base our decisions on our views about probabilities. But the costs and benefits, the value and disvalue of outcomes also plays a role in our decisions. For example, I am thinking about driving to Oklahoma City to see a movie, but Gary England said that there was a 40% chance of snow tonight, and my car doesn't handle well on slick roads. Should I go or not? If I don't want to see the show very badly I may stay home, but if this is my only chance to see something I've really been wanting to see, the trip may be worth the risk.

Odds of 2 to 1 may be enough for me to bet a few dollars, but not to bet my life (as I would in a case of risky surgery). Both the probabilities and the values (and disvalues) of outcomes play quite a role in our decisions. The following examples should help us see how this should work if we are reasoning well.

*Example 1: Three Point Shots* Wilma, one of the guards on the UCLA basketball team, hits 40% of her shots from less than three point range and 30% of her shots from three point range. It may be best for Wilma to take certain shots in certain cases (e.g., if two points will win the game then she should go for two). But in general is it better for her to take two point shots or three point shots? The probability of hitting a three pointer is lower, but the payoff is higher. How do we weigh these two considerations?

The following table gives us the answer:

	<b>Probability</b>	×	<b>Payoff</b>	=	<b>Expected Value</b>
<b>Two pointer:</b>	.40	×	2 points	=	.8 points
<b>Three pointer:</b>	.30	×	3 points	=	.9 points

Over the long haul Wilma will, *on average*, get 0.8 points for each two point shot she takes and 0.9 points for each three point shot. We say that .8 is the *expected value* of Wilma's two point shot and .9 is the expected value of her three point shots. Over the course of a season this difference can matter, and other things being equal it is better for Wilma to attempt three pointers.

*Example 2: Rolling Dice* Your friend asks you to play the following game. You roll a die. If you get a six, he pays you six dollars. If you don't get a six, you pay him one dollar. Would this be a profitable game for you to play? To answer this question, we need to determine the expected value of this game.

Expected value of shooting threes

The formula for this when two outcomes are possible is this:

$$\begin{array}{l} \text{Probability of success} \quad \times \quad \text{Payoff (positive)} \\ \text{Plus Probability of failure} \quad \times \quad \text{Payoff (negative)} \end{array}$$

In the case of two pointers and three pointers we could leave out probability of failure since the payoff in such cases is zero points. When we multiply this by the probability of failure, the result is still zero, so it drops out of the picture. But in the present case there is a “negative payoff” for failure.

Plugging the numbers in for the game proposed by your friend your expected value is determined by the following rule:

$$\begin{array}{l} \text{Probability of success (= } 1/6) \quad \times \quad \text{Payoff (= \$6)} \quad = 1 \text{ dollar} \\ \text{Plus Probability of failure (= } 5/6) \quad \times \quad \text{Payoff (= - \$1)} \quad = \underline{-5/6 \text{ dollar}} \\ \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad = 1/6 \text{ dollars.} \end{array}$$

The expected value of this game for you is  $1/6$  of a dollar. Over the long run your average winnings per roll will be  $1/6$  of a dollar, or about sixteen and a half cents. Over the short run this isn't much, but it could add up over time. So it's a good game for you (though not for your friend—unless he enjoys losing).

*Exercise:* What payoffs should your friend propose if he wants the game to be fair for both of you?

The treatment of expected values can be extended in a natural way to cover more than two alternatives at a time. Just list all of the possible outcomes, and record the probability and the payoff for each (listing losses as negative payoffs). Multiply the probability for each outcome by the payoff for that outcome. Then add up all of these numbers.

You should think a bit about expected value before you play the slot machines, buy tickets for a lottery, or the like. All of these cases there is a positive expected value for those running the game, a “house advantage,” and a negative expected value for those playing it. A similar point holds for insurance premiums. The insurance company calculates the probabilities of various outcomes and then determines prices of policies and amounts of payoffs so that the company will have a sufficiently high expected value for each policy.

There is a subjective side to payoffs. Even in games of chance, dollars aren't the only things that matter. Some people like gambling, and so even if they lose a little money over the long run their enjoyment compensates for this loss. Other people dislike risk, so even if they win a bit over the long run the overall value of the game is negative for them.

There are many other cases where payoffs involve a person's own feelings about things. Wilbur has a heart condition that severely limits the things he can do. The probability that a new form of surgery will improve his condition dramatically is about 50%, the chances he'll die in surgery are 7%, and the chances the surgery will leave him about the same are 43%. Should he get the surgery?

That depends on how much various things matter to Wilbur. If being alive, even in a very unpleasant physical condition, is really important to him, then his assessment of the payoffs probably means that he shouldn't elect surgery. But if he can't stand being bed-ridden, he may assess the payoffs differently.

### 16.2.1 Pascal's Wager

Blaise Pascal (1623–1662) was one of the founders of probability theory. He was also a devout Catholic in seventeenth-century France. He argued that we should believe in God for the following reasons. As long as we are on this earth, we can never really settle the matter of whether God exists or not. But either He does or he doesn't.

*Case one: God exists*

1. If God exists and I believe that He exists, then I get a very high payoff (eternal bliss).
2. if God exists and I do not believe in Him I get a very negative payoff (fire and brimstone for all eternity).

*Case two: God does not exist*

3. If God does not exist and I believe that He does, I make a mistake, but its consequences aren't very serious.
4. If He doesn't exist and I don't believe in Him, I am right, but being right about this doesn't gain me a lot.

Pascal uses these claims to argue that we should believe in God. What are the relevant probabilities, payoffs, and expected values in each case? Fill in the details of his argument. What are the strengths and the weaknesses of the argument?

### Exercises

1. Edna hits 45% of her three point shots and 55% of her two point shots. Which shot should she be trying for?

2. Suppose your friend Wilma offers to play the following game with you. You are going to roll a pair of dice. If you get a 7 or 11 (a natural) she pays you \$3. If you roll anything else, you pay her \$15. What is the expected value of the game for you? What is it for her?
3. Wilbur and Wilma are on their first date and have gone to the carnival. Wilma is trying to impress Wilbur by winning a stuffed toy for him. Wilma is trying to decide between two games: the duck shoot and the ring toss. She can shoot 55% of the ducks, which are worth two tickets each, and she can make about 35% of the ring tosses, which are worth four tickets each. Assuming that Wilma needs to accumulate 15 tickets to win the toy, which game should she play?
4. In an earlier chapter we learned about roulette. Calculate the expected value for betting on the number 13 (recall that the true odds against this are 37 to 1 but the house odds are 35 to 1).

### 16.3 The Gambler's Fallacy

We commit the **gambler's fallacy** when we treat things that are independent as though they were not independent, that is to say, when we (mistakenly) think that one of two independent things somehow influences or affects the other. For example, the outcomes of successive flips of a *fair* coin are independent of each other, so the outcome of the second flip does not depend in the least on the outcome of previous flips. If you flip a fair coin ten times and it comes up heads each time, the probability of it coming up heads on the eleventh flip is still  $1/2$ .

**Gambler's fallacy:** treating independent events as if they were not independent

Of course if you get enough heads in a row you may begin (quite reasonably) to suspect that the coin really isn't fair. But even if it is biased, so that it is likely to come up heads twice as often as tails, the point remains: the outcomes of the two successive flips are independent of each other, so what happens on the next flip isn't affected by earlier outcomes.

In such situations we have a tendency to think that the coin is more likely to come up tails in order to "even things out," to satisfy the "law of averages." But the coin doesn't "remember" what it did on earlier flips, and people who reason this way commit the gambler's fallacy. Similarly defective thinking is common with other games of chance like roulette, and it is a danger in any reasoning involving probabilities.

The gambler's fallacy is not restricted to games of chance. Suppose that Wilbur and Wilma have four children, all boys. They would like to have a girl, and they reason as follows. Very nearly half of the children born in the world are girls. We have had four boys in a row, so it's got to be time that we get a girl. It is an

*empirical question* whether having a child of one sex affects the probability of the sex of subsequent children. The evidence strongly suggests that it does not; the sex of one child is independent of the sex of its siblings. So, assuming that the gender of a couple's children are independent of one another, Wilbur and Wilma commit the gambler's fallacy.

There is a saying that lightning never strikes in the same place twice, and some people will even seek refuge in spot where lightning struck before in hopes of being safe. It may be true that lightning rarely strikes in the same place twice, but that is simply because the probability of it striking *in any specific spot* is reasonably low. But the lightning now doesn't know where lightning has struck before, and this general slogan "never in the same place twice" rests on the gambler's fallacy.

Nothing in these cases requires us to have any precise ideas about probability values. As long as we have good reason to think that two things are independent, we shouldn't act as though one could influence another. For example, we may have reason to think that Wilbur's die is loaded so that sixes are more likely to come up than any other number. We may not know how much more likely sixes are, but as long as the outcomes of separate throws are independent, only bad reasoning can lead us to suppose that since a six hasn't come up the last ten throws, a six must be due on the next throw.

## 16.4 The Conjunction Fallacy

We begin this section with two puzzles.

1. Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations. Which is more likely:

- Linda is a bank teller
- Linda is a bank teller who is active in the feminist movement

2. Which alternative seems more likely to occur within the next ten years?

- An all-out nuclear war between the United States and Russia
- An all-out nuclear war between the United States and Russia in which neither country intends to use nuclear weapons, but both sides are drawn into the conflict by the actions of a conflict that spirals out of control in the Middle East.

In both cases the second alternative is a conjunction that includes the first alternative as one of its conjuncts. So for the second option to be right (in either case), it must be possible for a conjunction to be more probable than one of its conjuncts. But this can *never* happen. But these examples attest to our tendency to judge some conjunctions more probable than their conjuncts. Since this involves bad reasoning, we will call it **conjunction fallacy**.

There are three ways to see that such reasoning is fallacious. First, we can think about what would have to be the case if it were right. How could the probability of two things happening together be greater than the probability of either one happening by itself? After all, for both to occur together, each of the two must occur.

Second, the point is a consequence of the rule for probabilities of conjunctions. This is easiest to see in the case where the conjuncts are independent, though the same idea applies in cases where they are not. Unless one of the conjuncts has a probability of 1, we will be multiplying a number less or equal to 1 by a number less than 1, and the result will have to be smaller than either of the two original numbers.

When  $A$  and  $B$  are independent,  $\Pr(A \& B) = \Pr(A) \times \Pr(B)$ . If the probability of each conjunct is 1, then the probability of the conjunction itself will be 1. But in most real-life situations, the probabilities of the two conjuncts is less than 1. In that case the conjunction will be less probable than either conjunct. For example, if  $\Pr(A) = .9$  and  $\Pr(B) = .9$ , the probability of the entire conjunction is only .81. If  $\Pr(A) = .7$  and  $\Pr(B) = .6$ , the probability of the conjunction is .42.

Third, and best, we can draw a diagram to represent the situation.




Figure 16.1: Feminist Bank Tellers

The crosshatched area where the circles overlap represents the set of bank tellers who are also feminists. Clearly this area cannot be larger than the entire

**Conjunction fallacy:** thinking a conjunction is more probable than either of its conjuncts

circle on the left which represents bank tellers.

### Specificity and Probability

More detail means  
lower probability

As we add detail to a description it often becomes more *specific*. And as it becomes more detailed and specific, it becomes less probable. For example, suppose that you are going to toss a quarter once. The probability of its landing heads is  $1/2$ . But the probability it will land heads with Washington looking more or less north is less, and the probability he'll be looking due north is very small. Indeed, the probability he'll be looking in *any* direction that you specify precisely before the flip is minuscule.

This bears directly on the conjunction fallacy, because adding more detail is really just a matter of adding more conjuncts, and adding more conjuncts typically adds more detail. To say that the quarter will land heads with Washington looking north is to say that it will land heads *and* Washington will be looking north. And as always, a conjunction cannot be more probable than either of its conjuncts.

Finally, note that we don't need to know precise probability numbers to appreciate the fundamental point that the probability of the conjunction can *never* be greater than the probability of its least probable conjunct, whatever its probability might turn out to be.

## 16.5 Doing Better by Using Frequencies


Think Frequencies!

A good deal of research has shown that we reason more accurately about many probabilities, including the probabilities of conjunctions, if we think in terms of *frequencies* or *proportions* or *percentages*, rather than simply in terms of probabilities. Recall Linda, the single, outspoken, bright, philosophy major. When people are asked whether it is more probable (or more likely) that she is (1) a bank teller, or (2) a bank teller who is active in the feminist movement, well over half of them usually (incorrectly) select (2).

But when people approach the same problem in terms of percentages or frequencies, they do better. If we keep the same profile, but rephrase the two questions to ask: what proportion or percentage of a group of one-hundred randomly selected women who fit this profile are (1) bank tellers, and what proportion or percentage of a group of one-hundred randomly selected women who fit this profile are bank tellers who are active in the feminist movement, more people avoid the conjunction fallacy. More people (correctly) select (1)—although there is still a strong tendency to commit the conjunction fallacy.



This tendency also shows up in frequency versions of the conjunction fallacy. Are there are more six letter words ending in ‘ing’ than having ‘n’ as their fifth letter is a question about relative frequencies. Many still say that there are more with ‘ing’, even though every ‘ing’ ending as an ‘n’ as its fifth letter, and there are also *non-‘ing’* words with ‘n’ in fifth place (e.g., ‘barons’). Still, many of us do *better* here if we think in terms of percentages or proportions or frequencies than if we simply think in terms of probabilities than if we simply think in terms of probabilities.



In short, one of the best ways to improve your accuracy in estimating probabilities is to rephrase things in terms of frequencies whenever you can. Instead of asking how probable it is that a person with a given set of symptoms has a particular disease, ask “what *proportion* of people in a randomly selected group of 100 who have these symptoms have this disease?” What is the *frequency* of this disease in a group of 100 people who have these symptoms. In fact, you don’t even need words like ‘frequency’ or ‘percentage’. Just ask: about how many people out of a hundred (or a thousand) who have these symptoms also have the disease. You can then translate your answer into percentages or probabilities very easily.

It may also help to use percentages instead of probability numbers. Rather than saying the probability of having the disease is .2 you can say that the probability is 20%. When you ask how many things out of a hundred have a certain property and use percentages, then the percentages translate directly into number of things; 90% is just 90 of the items out of 100.

Thinking in terms of percentages or frequencies also makes it easier to think about cumulative risk. If the probability of a particular brand of condom failing is 0.01 ask: how many times out of 100, or 1000, would it fail. The respective answers are 1 every 100 times, and 10 every thousand times. So over the longer run there is a substantial chance of failure.

## 16.6 Why Things go Wrong

For a spacecraft to make it to the surface of Mars many different subsystems have to function properly. The computer, the radio, the rocket engines, and more must all work. If any of them fail, the entire mission may be ruined. In the case of the Challenger, a problem with the O-rings was enough for catastrophic results.

Similar points apply to many other cases. For a computer, or a car or a computer program, to work all of its parts must work. For the human body to remain healthy, all of the various “systems” need to work; the cardio-vascular system, the

Success often requires a conjunction

immune system, the nervous system, and many other things must each function well if we are to remain healthy. In general, a complicated system will only function properly if a number of its parts each functions properly.

To say that all of the subsystems of something must work is to say that subsystem one must work *and* subsystem two must work *and* subsystem three must work, and so on. Your computer's central processor, and its hard drive, and its monitor, and . . . all have to work for the computer to work. This means that a conjunction must be true. And as we saw in the previous section, the probability of a conjunction is usually lower than the probabilities of its conjuncts.

Imagine a spacecraft consisting of five subsystems. If any one of them fails, the entire mission will fail. Suppose that the probability that each subsystem will work is .9 and that the performance of each is independent of the performance of the others (this simplifying assumption won't actually be the case, but it doesn't affect the present point in any relevant way). Then the probability that all five of the subsystems will work is  $.9 \times .9 \times .9 \times .9 \times .9 (= .9^5)$ , which is a bit less than .6. If there were seven subsystems (all with a .9 probability of working), the probability that all seven would function properly is less than .5.

Even if the probability that each part of a complex system will function correctly is .99, if there are enough systems failure somewhere along the line is likely; the more components there are, the more the odds against success mount up (this is why spacecraft typically include backup systems).

Failure often only requires a disjunction

We can make the same point in terms of disjunctions. A chain is no stronger than its weakest link; if any of them break, the entire chain gives way. If the first one breaks, *or* the second, *or* the third, . . . , the chain is broken. Many things are like chains; they can breakdown in several different ways. In many cases the failure of one part will lead to the failure of the whole.

In our imaginary spacecraft, the failure of subsystem one *or* of subsystem two *or* of subsystem three *or* . . . can undermine the entire system. This is a disjunction, and the probabilities of disjunctions are often larger than the probabilities of any of their disjuncts (this is so because we *add* probabilities in the case of disjunctions).

The lessons in this section also apply to things that occur repeatedly over time. Even if the probability of something's malfunctioning on any particular occasion is low, the cumulative probability of failure over a long stretch of time can be moderate or even high.

Contraceptives are an example of this. On any given occasion a contraceptive device may be very likely to work. But suppose that it fails (on average) one time out of very 250. If you use it long enough, there is a good chance that it will eventually let you down. To take another example, the chances of being killed in an automobile accident on any particular trip are low, but with countless trips over the years, the odds of a wreck mount up.

There is considerable evidence that people tend to overestimate the probabilities of conjunctions (thinking them more likely than they really are) whereas they underestimate the probabilities of disjunctions (thinking them less likely than they really are). As a result, we tend to overestimate likelihood of various successes while we underestimate the likelihood of various failures.

As before, you don't need to know precise probability values to appreciate these points. You may know that a contraceptive is pretty likely to work but that there is a non-negligible chance it will fail. This tells you that over time there is a very real chance of its failing.

## 16.7 Regression to the Mean

The grade school in Belleville administers an achievement test to all the children who enter fifth grade. At the end of the school year they give the same test again. The average score both times is 100, but something odd seems to have happened. Children who scored below average on the test the first time tend to improve (by about five points), and children who scored above average tend to do worse (by about five points).

What's going on? Might it be that when the two groups of children interact extensively, as they did over the school year, the higher group pulls the other group up while the lower group pulls the higher group down?

Here is another example. Instructors in an Israeli flight school arrived at the conclusion that praising students for doing unusually well often led to a decline in their performance, while expressing unhappiness when they did poorly often led to an improvement. This group just happened to be studied, but many other instructors and teachers come to similar conclusions. Are they right?

The more likely explanation in both cases is that they involve regression to the mean. **Regression to the mean** is a phenomena where more extreme scores or performances tend to be followed by more average ones. It's also called **mean reversion**.

The basic idea is that extreme performances tend to revert or regress or move back toward the mean (i.e., toward the average). So unusually low scores tend to be followed by higher ones (since low scores are below the mean). And unusually high scores tend to be followed by lower ones (since high scores are above the mean). Regression to the mean can occur with anything that involves chance, it occurs frequently, and it is very easy to overlook it.

Suppose that you did unusually well (or unusually badly) when you took the ACTs. There is a reasonable probability that if you took them again your second score would be closer to the average. In sports someone who has an unusually

**Regression to the mean:**  
more extreme performances  
tend to be followed by more  
average ones

good or unusually bad game is likely to turn in a more average performance the next time around. Indeed, people often remark on the sophomore slump, in which athletes who did exceptionally well as freshmen fall off a bit as sophomores, and the *Sports Illustrated* jinx, in which people who do very well and make it to the cover of *Sports Illustrated* play worse in subsequent weeks. Many cases of this sort simply involve regression to the mean.

Since regression to the mean can occur with anything that involves chance, it affects more than just performances on the basketball court or in the concert hall. For example, there is a lot of randomness in which genes parents pass on to their offspring, and parents who are extreme along some dimension (unusually tall or short, unusually susceptible to disease, etc.) will tend to produce children whose height or susceptibility to disease are nearer the average. Two very tall parents are likely to have tall offspring, but the children are not likely to be as tall as the parents (or not as tall as the father, in the case of boys).

Why do things tend to “regress” to the average value? Why not to some other point? Performances involve a “true level of ability” plus chance variation (“error”). The chance variation can involve many different things that lead to better or worse scores than we would otherwise have.

Suppose that you take the ACT several times. Some days you may be very tired, other days well rested; some days nervous, other days more focused and confident; some days you may make a lot of lucky guesses, other days mostly unlucky ones. Often the good and bad conditions will pretty much cancel out, but sometimes you will have mostly the good conditions (in which case you will score very well), and sometimes mostly the bad ones (in which case you will score poorly). If you score extremely well, the chances are that this is a combination of a high ability plus auspicious background conditions, and so your score is likely to be lower the next time around.

Because of the chance error, the distribution of your performances fits a pattern that resembles the standard bell curve. In this distribution of scores, the average value is the value closest to the largest number of cases (we will return to this point when we consider descriptive statistics). So unusually good or unusually bad performances are likely to be followed by the more probable performances, which are just those nearer the average.

The idea may be clearer if we consider a concrete example. Suppose that you shoot thirty free throws each day. Over the course of a month the percentage of shots that you hit will vary. There is some statistical variation, “good days” and “bad days”.

Many things may improve or weaken your performance: how sore your muscles are, how much sleep you got, how focused you are. Sometimes all the things come together in the right way and you do unusually well; other times everything

seems to go wrong. But on most days these factors tend to cancel each other out, and your performance is nearer your average. Since your performance is more often near the mean, extreme performances are likely to be followed by more average ones.

### 16.7.1 Regression and Reasoning

Regression to the mean is very common but frequently overlooked, and failure to appreciate the phenomenon leads to a lot of bad reasoning.

#### Regression and Prediction

Suppose that Wilbur has an exceptionally good or an exceptionally poor performance shooting free throws in a game. It is natural to base our prediction about how he will do next time on his free throw percentage in the game we saw; we just project the same percentage. But if his shooting was way above (or way below) the average for players in general, his percentage in subsequent games is likely to regress toward this average.

Again, a company may detect falling profits over the previous three months. The manager gets worried, thinks about a way to change marketing tactics, and predicts that this will turn things around. The new tactics are adopted and profits go back up to their previous level. But this may result simply from regression to the mean. If so, the new marketing techniques will (incorrectly) be given credit for the turnaround.

There are various cases where our failure to take regression into account leads to bad predictions; for example, if someone does unusually well (or unusually poorly) in a job interview, we are likely to have a skewed impression of how well they will do on the job.

#### Explanation and Superstition

Suppose that you've had a couple of poor performances of late. Things went badly on some exams or in the last two recitals you gave. Then you have an unusually good performance the day you wear your green sweater, the ugly one your aunt gave you, and it may become your lucky sweater. Superstitions are often based on the fact that something (e.g., wearing the lucky sweater) just happened to coincide with a shift toward a better performance that is simply due to regression to the mean.

Of course most of us don't really believe in lucky sweaters (though we might still wear one, on the theory that "it can't really hurt"). But lack of awareness

of regression to the mean is responsible for a lot of bad reasoning. Whenever an element of chance is involved, regression to the mean comes into play. Nisbett and Ross note that if there is a sudden increase in something bad (e.g., an increase in crime, divorce rates, bankruptcies) or sudden decrease in something good (e.g., a decline in high-school graduation rates, in the amount given to charity) some measure is likely to be taken. For example, if there is a sudden increase in crime, the police chief may increase the number of police officers walking the beat.

If the implementation of a new policy is followed by a decrease in something undesirable or an increase in something desirable, we are likely to conclude that the measure is responsible for the shift. But in many cases such a shift would have occurred without the measure, simply as a consequence of regression to the mean. In such cases, we are likely to *explain* the drop in crime by the increased number of police on the beat. The measure will be given too much credit.

As a final example, let's return to the question of rewards and punishments. Parents and teachers often have to decide whether rewards or punishments are more likely to be effective. Unusually good behavior is likely to be followed by less good behavior simply because of regression, and unusually bad behavior is likely to be followed by better behavior for the same reason.

Hence, when we reward someone for doing extremely well they are likely not to do as well the next time (simply because of regression to the mean). Similarly, if we punish them for doing badly they are likely to do better next time (for the same reason). In each case the change in performance may simply be due to regression to the mean, and the reward and punishment may have little to do with it. It will be natural to assume, though, that punishments are more effective than rewards.

When something like punishment or an increase in police on the beat accompanies regression to the mean we can easily conclude that society in general, or that we in particular, have found a method to solve certain sorts of problems when in fact we have little power to solve them. Obviously this doesn't make for good decision making at either the public or the personal level.

## 16.8 Coincidence

Some things strike us as very unusual and unlikely. This often leads us to think that there "must be something special going on" when they do occur—surely they couldn't "just happen by chance." Wilbur survives a disease that is fatal to 99.8% of the people who contract it, so something special *must* be going on. In fact, though, there will be two people out of every thousand who do survive. When Wilbur's doctor first saw the test results, he thought it very unlikely that Wilbur would make it, and when Wilbur does the doctor is amazed. But there must be two people who

are the lucky pair in a thousand, and it may just have happened to be Wilbur.

As we saw earlier, if we describe an event in enough detail, it will seem very unlikely (before the fact) that it will occur. Suppose you toss a quarter ten times. The probability of *any* particular sequence of outcomes is  $(1/2)^{10}$ , which means that each possible set of outcomes is extremely unlikely. But when you actually do the tossing, *one* of these very unlikely sequences will be the one that you actually get.

There are countless examples of this. Before you tee off, the probability of the ball landing in any particular spot is close to zero. But if you hit it, it *will* land some place or another, even though it was very unlikely that it would alight precisely where it eventually does. So things that seem unlikely can, and do, happen just by chance. Indeed, if we describe things in enough detail, almost everything that happens would have seemed unlikely before it occurred. Still, one of these unlikely things will occur.<sup>3</sup>

## 16.9 Chapter Exercises

1. Evaluate the following argument in light of concepts recently covered in class.

The burglary rate here in Belleville has always been very close to 1 robbery per 400 homes. But last year it ballooned up to 3.9 per 400. However, the Chief of Police quickly hired three additional policemen, and this year the burglary rate is back down to where it had been before last year's increase. So hiring more police is a good way to lower the burglary rate.

2. Children doing below average work in school who suddenly do well on an achievement test are often labeled underachievers. Sometimes they are, but what else might be going on?
3. I often find that I am disappointed when I return to a restaurant that seemed outstanding on my first visit. I'm often tempted to conclude that the chefs got

---

<sup>3</sup>The example of Israeli flight instructors and related cases of regression to the mean is discussed by Daniel Kahneman & Amos Tversky in "On the Psychology of Prediction," *Psychological Review*, 80 (1873); 237–251. Many related papers by Tversky, Kahneman and others are reprinted in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under Uncertainty: Heuristics and Biases*, Cambridge University Press, 1982. Richard E. Nisbett and Lee Ross's discussion of regression to the mean occurs in ch. 7 of their *Human Inference: Strategies and Shortcomings of Social Judgment*, Englewood Cliffs, NJ: Prentice-Hall, 1980. John Paulos' *Innumeracy: Mathematical Illiteracy and its Consequences* contains excellent and very accessible discussions of chance, magnitudes of numbers, and coincidence.



lazy over time or that the management quite working as hard as they had at the beginning. What do you think about my reasoning? What other explanations might there be for this result?

4. Wilbur reasons in the following way: If a method of contraception has a 6% failure rate, then we should expect the same probability of getting pregnant in 1 year of use as we would in 10 years of use. The chances are that 6% of the people using it will create a pregnancy.

5. If a student gets the highest grade in her class on the first examination in Critical Reasoning what grade would you predict that she will get on the midterm? Justify your answer?

6. Most of you had to take the ACT or the SAT test; if you apply to graduate school you will have to take the GRE, and if you apply to Law School you will have to take the LSAT. These are multiple-choice tests. To discourage random guessing many tests of this sort subtract points for wrong answers. Suppose that a correct answer is worth +1 point and that an incorrect answer on a question with 5 listed answers (a through e) is worth  $-1/4$  point.

1. Find the expected value of a random guess.
2. Find the expected value of eliminating one answer and guessing between the remaining 4 possible answers.
3. Using your answers to (a) and (b), when would it be advisable to guess and why?

7. Suppose that you take a multiple choice exam consisting of ten questions. Each questions has four possible answers. The topic is one you know nothing about and you are reduced to guessing. What is the probability that you will guess right on the first question? What is the probability that all ten of your guesses will be right (you know nothing about the subject matter, so you guess at random and so you can assume independence)? If one million people took the exam how good are the chances that at least one person would get all of the answers right simply by random guessing (don't worry about assigning a number to this, but be as precise as you can and justify your answer).

8. From a lecture on fire safety in the home: "One in ten Americans will experience some type of destructive fire this year. Now, I know that some of you can say that you have lived in your home for twenty-five years and never had any type of fire. To that I would respond that you have been lucky. . . . But that only means that you are not moving farther away from a fire, but closer to one." Evaluate the reasoning in this passage.



9. Joe makes an average of 35% of his basketball shots. After playing in a pick-up game in which he misses all six of the shots he takes, he argues that in the next game he will be hot because, having missed six already, he has the odds in his favor. Evaluate Joe's reasoning.

10. Suppose that you built a computer that had 500 independent parts. And suppose that each part was 99% reliable when used the very first time. What are the chances that such a computer would work the very first time it was turned on?

11. In the previous chapter we noted Laurie Anderson's quip: The chances of there being two bombs on a plane are very small, So when I fly I always take along a bomb. We are now in a better position to analyze the bad reasoning involved. Do so.

#### Answers to Selected Exercises

11. The chances of there being two bombs on a plane are very small, So when I fly I always take along a bomb. Assuming that I am not in league with any terrorists, whether I bring a bomb has no effect on whether someone else also brings a bomb along on the flight. The two events are independent. But the joke treats them as though they were dependent (my bringing a bomb makes it less likely that others will). Hence it involves a subtle instance of the gambler's fallacy.



## **Part VII**

# **Systematic Biases and Distortions in Reasoning**



## **Part VII. Systematic Biases and Distortions in Reasoning**

In this part we will examine several common errors and biases in our everyday reasoning.

In Chapter 17 we will see how we often rely on rough-and-ready strategies called *heuristics* in our reasoning. In many cases heuristics allow us to draw reliable inferences quickly. But if we rely on them too heavily, or in the wrong situations, they can lead to bad reasoning.

In Chapter 18 we will study several further biases in our thinking.

In Chapter 19 we study peoples' perceptions of their own inconsistency and the ways these influence their attitudes, action, and thought.



# Chapter 17

## Heuristics and Biases

**Overview:** We often rely on rough-and-ready strategies called *heuristics* in our reasoning about things. In many cases these heuristics allow us to draw reliable inferences quickly. But if we rely on them too heavily, or in the wrong situations, they lead to bad reasoning. In this chapter we will study several inferential heuristics, examine the ways they promote faulty reasoning, and devise safeguards against such errors.

### Contents

---

<b>17.1 Inferential Heuristics</b> . . . . .	<b>330</b>
17.1.1 Sampling Revisited . . . . .	330
<b>17.2 The Availability Heuristic</b> . . . . .	<b>331</b>
17.2.1 Why Things Are Available . . . . .	332
<b>17.3 The Representativeness Heuristic</b> . . . . .	<b>335</b>
17.3.1 Specificity Revisited . . . . .	337
<b>17.4 Base-Rates</b> . . . . .	<b>337</b>
<b>17.5 Anchoring and Adjustment</b> . . . . .	<b>340</b>
17.5.1 Anchoring Effects can be Very Strong . . . . .	341
17.5.2 Anchoring and Adjustment in the Real World . . . . .	341
17.5.3 Safeguards . . . . .	342
<b>17.6 Chapter Exercises</b> . . . . .	<b>342</b>

---

## 17.1 Inferential Heuristics

Human beings have many limitations. We have limited memories, attention spans, and computational abilities. We also have better things to do than to spend our time trying to reason precisely about everything that we ever think about. So we use shortcuts. These shortcuts are called inferential (or judgmental) heuristics.

An inferential (or judgmental) **heuristic** is a general strategy that we use for drawing inferences. It is a rough-and-ready device, a cognitive shortcut, a rule of thumb for reasoning. We use inferential heuristics frequently, usually without being aware of it. Heuristics are usefully contrasted with definite and specific rules for reasoning (like our earlier rules for calculating probabilities)

Like many of the cognitive mechanisms studied in earlier chapters, inferential heuristics are often quite useful. They allow us to draw rapid inferences without having to gather data or compute probabilities. This is valuable, because we rarely have the time, energy, know-how, or interest to go to such trouble. Indeed, in order to survive organisms must be able to process information and draw conclusions quickly, and handy, habitual rules of thumb—heuristics—are often better suited for this than more reliable, but time-consuming, rules for reasoning. The drawback is that overreliance on inferential heuristics can lead to serious biases or errors in reasoning.

### 17.1.1 Sampling Revisited

**Good sample:**

1. big enough
2. representative

Remember how inferences from samples to populations work. When we infer a conclusion about a population from a description of a sample from it:

1. the *premises* are claims about the *sample*.
2. the *conclusion* is a claim about the *population*.

For example, we might draw a conclusion about the average income of Oklahomans based on the results of a sample of 2000 Oklahomans.

Our conclusion involves an inductive leap. It goes beyond the information in its premises, because it contains information about the *entire population* while the premises only contain information about the *sample*. But if we satisfy two conditions involving the sample, our inference can still be inductively strong. We must have

1. a large enough sample
2. a representative (unbiased) sample

An unbiased sample is typical of the population. By contrast, in a biased sample, some portions of the population are overrepresented and others are underrepresented.



The problem with a very small sample is that it is unlikely to be representative. Other things being equal, a bigger sample will be more representative. But there are costs to gathering information, costs in time and dollars and energy, so it is rarely feasible or desirable to get samples that are huge.

## 17.2 The Availability Heuristic

- Do more people in the U.S. die from murder or suicide?
- Are there more English words that begin with the letter ‘r’ or words that have ‘r’ as their third letter?
- Are there more famous people from Oklahoma or from Kansas (a state with roughly the same population)?

Many of our inferences lead to conclusions about the *relative frequency* or proportion of some feature in a given population. Are there more *Bs* than *Cs*; e.g., are there more Jeeps (*Bs*) or Fords (*Cs*)? Is an *A* more likely to be a *B* or a *C*; e.g., is an OU student (an *A*) more likely to be male (*B*) or female (*C*)? Since we can rarely check the entire population, we must base our inference on a sample. In everyday life we rarely have the time or resources to gather a sample in the way scientists do, so we often “do the sampling in our heads.” We try to remember cases that we know about or to imagine cases that seem relevant.

Suppose that you want to know whether there are more Fords or Jeeps in use today. You will probably rely on the sample of vehicles you can recall. You try to think about the various makes of vehicles you have observed. Obviously this method is somewhat vague and impressionistic, since you probably don’t remember more than a handful of specific Fords (like Wilbur’s Crown Victoria) or Jeeps (like Aunt Ethel’s). But at least you do know that you have seen a lot more Fords than Jeeps. You have a generalized memory about this, even though you don’t recall many specific vehicles of either kind.

In many cases, including this one, this method works. You remember seeing a lot of Fords, and you remember that you haven’t seen many Jeeps. The reason why you remember more Fords is that there *are* more Fords. Fords are easily available to recall largely because there are so many of them. Fords are more **available** in your memory precisely because you have seen a lot more Fords.

When we need to judge the relative frequency or probability of something, we are often influenced by the **availability** or accessibility of those kinds of things in thought. The “sample in our heads” consists of the cases we remember and, to some extent, the cases we can easily imagine. We use the **availability heuristic** when we base our estimates of frequencies or probabilities on those cases that most readily

**Availability heuristic:** basing judgments of frequency on ease with which examples can be recalled

come to mind, on those that are most *available* in memory and imagination. This heuristic inclines us to assume that things that are easier to remember or imagine are more likely to occur.

### 17.2.1 Why Things Are Available

Often we remember certain things because they really do occur frequently, and when this is the case the available sample in our head will often be a good one (or at least good enough for the rough-and-ready inferences of everyday life). When availability is highly correlated with objective frequencies or probabilities, as it often is, it is a useful guide. Are there more words beginning with the letter 'r' or with the letter 'z'? Words beginning with 'r' are more available than those beginning with 'z' precisely because they are much more common. Here the heuristic works very well, leading us to the correct conclusion.

But the *judged frequency* and the *true frequency* of something may be very different. Things may be available in memory or imagination for reasons having little to do with their frequency or probability. In these cases the availability heuristic leads us to rely on a small sample (the cases we easily remember) and one that may be biased in various ways (the cases we happen to have encountered and manage to recall). For example, things we are familiar with will be available. And since memory generally becomes less vivid and accessible over time, more recent experiences and events are more likely to be available than those that occurred longer ago.

We begin this section by asking whether there are more English words that begin with the letter 'r' or words that have 'r' as their third letter. Many of us think (at least when we aren't primed to think there is a trick involved) that more words begin with 'r', though in fact this is false. So why do we think this? But it is much easier to think of, to generate, letters that begin with 'r' than to think of words in which 'r' comes third. Words beginning with 'r' are *more available*, and this greater availability leads many of us to infer that more words begin with 'r'.

#### Examples of Availability

Fewer than one in a million Americans are killed by terrorists, whereas over one in 5000 are killed in automobile accidents. But the cases where Americans *are* killed by terrorists are likely to make the news, and for obvious reasons they stand out in memory. Now suppose we are asked to estimate the number of people killed by terrorists. The sample that is readily available, the one that comes naturally to mind, can easily lead us to overestimate the threat of terrorism today. People also overestimate the rate of homicides and other things that make the news. This is one

reason why most of us suppose that there are more murders than suicides, though statistics show that there are many more suicides than murders. By contrast, the frequency of things that are not so well-publicized, like death from diabetes, is usually radically underestimated.

On the other hand, unless we know of several people who have been killed in automobile accidents, examples of such deaths may not be so salient in memory. Such deaths are common enough that they aren't likely to be reported by the media unless the person killed is well-known. Examples of such deaths are not particularly available, so we may radically underestimate their frequency. To take a related example, fires make the news more often than drownings, and they may be more dramatic in various ways. So it is not surprising that many people think that death by fire is more likely than drowning, even though the reverse is the case. The good news here is that we may underestimate the amount of helping and kindness there is, since such things rarely make the news.

Things that occur reasonably often (e.g., fatal automobile accidents) are rarely reported and are easily forgotten, whereas things that are rare but dramatic (e.g., terrorism) make for good news and stand out in memory. In such cases, frequency is *not* closely related to availability in memory, and the use of the availability heuristic will lead us astray. For example, 100 times as many people die from disease as are victims of homicide, but newspapers carry three times as many articles about murders.

### Media Effects

Here are some further examples. The media and advertisers often tell us about people who struck it rich by winning a state lottery. This can make such cases more available to us in thought, leading us to overestimate the probability of winning a lottery (we all know the probability of winning is low, but it is *much* lower than many people suppose). Again, many more people die of diabetes each year than in accidents with fireworks. The latter get more press, however, and many people think more deaths really are caused by such accidents.

Partly because they are reported and partly because of the success of the movie *Jaws*, shark attacks seem vivid, easy to imagine, and easy to remember. In fact they are rare, and you are much more likely to be killed in many other (less dramatic) ways. A student of mine once told me that your chances of being killed by a pig are considerably higher than your chances of being killed by a shark. For many years this seemed plausible, but it was only when I sent students out on the internet to check that I learned that it is.

### Surprising Events are Memorable

There are many other cases where unlikely events may be particularly available. A few people are very likely to recover from an illness that is fatal to the vast majority of those who contract it. Since the tiny minority who recover will probably be under some sort of treatment (call it treatment *X*), and since miracle recoveries make for good news, we may hear about the miracle cure due to treatment *X*. This will be available to memory, and so we overestimate the probability that *X* can be effective in curing the disease.

Indeed, in all but the most extreme conditions *almost any* miracle cure or quick fix (for losing weight, kicking cigarettes, quitting gambling, etc.) will seem to work for some people (perhaps because of a placebo effect, perhaps thorough sheer coincidence). In such cases we may hear an endorsement, perhaps in an infomercial, from people who sincerely believe that they have benefited from the treatment. Such testimony can be very compelling, and it is often easily available in memory. In such cases the availability heuristic can lead us to spend a lot of money on quick fixes that don't fix anything at all (except the financial condition of the person selling them).

### Saliency

Aunt Ethel vs. *Consumer Reports*

One or two examples may be so vivid or salient that they lead us to discount much better evidence. Cases "close to home" can be especially compelling. Your Aunt Ethel had a Ford Crown Victoria that was a real piece of junk (though 'junk' wasn't exactly the word she used). This single case is likely to loom very large in your memory. Then you learn that some consumer group you trust (e.g., *Consumer Reports*) did a survey of thousands of car owners and found Crown Victorias to be more reliable than most other makes. If you are like most people, the one case close to home will stand out more (be more *salient*); it will be more memorable. Hence it will have a much greater influence on what you buy than the careful and detailed study by the consumer group.

### Our Everyday Samples are Often Biased

Many of the samples we encounter as we go about our lives are biased. Our age, gender, race, job, friends, interests, where we live all mean that we will be exposed more to some things than others. If I live in Boston, Massachusetts I will be exposed to a different range of things than if I live in Belleville, Kansas.

In many cases this is obvious, and it's relatively easy to discount for it. I realize that it's not safe to predict the general public's tastes in music on the basis of the musical tastes of the people I know; they don't provide a representative sample. But

in other cases the biased nature of the samples we normally encounter may be less obvious. It can be tempting, for example, to form beliefs about the general public's political views on various issues on the basis of the views we hear expressed most often. But these may not be representative of people's views in general.

### Problems with Availability

In earlier chapters we encountered several phenomena which suggest that the availability of things in memory is not always a good guide to how things really are. Perceptual set will incline us to notice certain things while overlooking others, thus influencing what makes it into memory in the first place. Then elaboration in memory can affect what we remember, as can the context in which we remember it. Further biases may enter because of primacy, recency, or halo effects. In short, the sample in our heads is often based on limited experience, and it can then be further distorted in a variety of ways.

Prejudices and stereotypes are an especially insidious example of this. If you have a negative stereotype of members of a certain group you are likely to notice some things (e.g., cases where a member of the group fails) than others (e.g., cases where a member succeeds). You will also be more likely to remember such cases, and find it easier to imagine them. When you then have to predict how typical members of that group will do, the negative cases will be more available than the positive ones, and you are likely to conclude that they will probably do poorly. We will return to this topic in a later chapter.

We can't abandon the availability heuristic. It is deeply ingrained in the way we reason, and it often works very well. But we need to be aware of the ways it can lead to fallacious reasoning. In particular, we need to realize that the samples in our heads (and in the heads of others, often even those in the heads of experts) are biased in one way or another.

## 17.3 The Representativeness Heuristic

Mike is six two, weighs over two hundred pounds (most of it muscle), lettered in two sports in college, and is highly aggressive. Which is more likely?

1. Mike is a pro football player.
2. Mike works in a bank.

Here we are given several details about Mike; the profile includes his size, build, record as an athlete, and aggressiveness. We are then asked about the *relative frequency* of people with this profile that are pro football players compared to those with the profile who are bankers.

What was your answer? There are almost certainly more bankers who fit the profile for the simple reason that there are so many more bankers than professional football players. We will return to this matter later in this chapter; the relevant point here is that Mike *seems* a lot more like our picture of a typical pro football player than like our typical picture of a banker. And this can lead us to conclude that he is more likely to *be* a pro football player.

Many of us make just this sort of error with Linda. Linda, you may recall, is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and she participated in antinuclear demonstrations. On the basis of this description, you were asked whether it is more likely that Linda is (i) a bank teller or (ii) a bank teller who is active in the feminist movement. Although the former is more likely, many people commit the conjunction fallacy and conclude that the latter is more probable.

What could lead to this mistake? Various things probably play some role, but a major part of the story seems to be this. The description of Linda fits our profile (or stereotype) of someone active in today's feminist movement. Linda strongly resembles (what we think of as) a *typical* or *representative* member of the movement. And because she *resembles* the typical or representative feminist, we think that she is very likely to *be* a feminist. Indeed, we may think this is so likely that we commit the conjunction fallacy.

We use the **representativeness heuristic** when we conclude that the more like a representative or typical member of a category something is, the more likely it is to be a member of that category. Put in slightly different words, the likelihood that  $x$  is an  $A$  depends on the degree to which  $x$  resembles your typical  $A$ . We reason like this:  $x$  seems a lot like your typical  $A$ ; therefore  $x$  probably is an  $A$ .

Sometimes this pattern of inference works, but it can also lead to very bad reasoning. For example, Linda resembles your typical feminist (or at least a *stereotype* of a typical feminist), so many of us conclude that she is likely to be a feminist. Mike resembles our picture of a pro football player, so many of us conclude that he probably is. The cases differ because with Linda we go on to make a judgment about the probability of a conjunction, but with both Linda and Mike we are misusing the representativeness heuristic.

Overreliance on the representativeness heuristic may be one of the reasons why we are tempted to commit the gambler's fallacy. You may believe that the outcomes of flips of a given coin are random; the outcomes of later flips aren't influenced by those of earlier flips. Then you are asked whether sequence HTH-HTHTT is more likely than HHHHTTTT. The first sequence may seem much more like our conception of a typical random outcome (one without any clear pattern), and so we conclude that it is more likely. Here the representative heuristic leads

us to judge things that strike us as representative or normal to be more likely than things that seem unusual.

### 17.3.1 Specificity Revisited

We have seen (p. 314) that the more detailed and specific a description of something is, the less likely that thing is to occur. The probability of a quarter's landing heads is  $1/2$ , the probability of its landing heads with Washington looking north is considerably less. But as a description becomes more specific, the thing described often becomes more concrete and easier to picture, and the added detail can make something *seem* more like our picture of a typical member of a given group.

Higher specificity means lower probability

In Linda's case we add the claim that she is active in the feminist movement to the simple claim that she is a bank teller. The resulting profile resembles our conception of a typical feminist activist, and this can lead us to assume that she probably *is* a feminist activist. This may make it *seem* likely that she is a feminist activist. And this in turn makes it seem more likely that she is a bank teller *and* a feminist activist than that she is just a bank teller. But the very detail we add actually makes our claim, the conjunction, less probable than the simple claim that Linda is a bank teller.

In short, if someone fits our profile (which may be just a crude stereotype) of the average or typical or representative kidnapper, spinster, or computer nerd, we are likely to weigh this fact more heavily than we should in estimating the probability that they are a kidnapper or a spinster or a computer nerd. This is fallacious, because in many cases there will be many people who fit the relevant profile who are not members of the group.

## 17.4 Base-Rates

A group of men in Belleville consists of 70 engineers and 30 lawyers. Suppose that we select Dick at random from the group. The following is true of Dick:

**Base rate fallacy:** ignoring or underutilizing base rates when estimating probabilities and making predictions

Dick is a 30 year old man, married, no children. He has high ability and high motivation, and promises to be quite successful in his field. He is well liked by his colleagues.

Based on this: Is Dick more likely to be an engineer, a lawyer, or are these equally likely? What's relevant to deciding?

Kahneman and Tversky told subjects that they were dealing with a pool of a hundred people, 70 of whom were engineers and 30 of whom were lawyers. If they were simply asked to estimate the likelihood that some person, Dick, selected at



random from this group was an engineer, most said 70%. Another group was given the above description of Dick. The important thing about this description is that it is an equally accurate description of a lawyer or an engineer (and most subjects in pretests thought so).

**Dilution effect:** tendency for irrelevant (nondiagnostic) information to dilute or weaken relevant (diagnostic) information

The information in the description could be of no help in estimating whether someone is a lawyer or engineer, so we should ignore it and (in the absence of any other relevant information) simply go on the base rates. This means that we should conclude that the probability that Dick is a lawyer is .7. In the absence of the irrelevant description, people did just this. But when they were given the *irrelevant* description, they concluded that the probability that Dick was a lawyer was .5 (fifty/fifty). The irrelevant information led them to disregard base rates; they simply threw away information that is clearly relevant.

This is an instance of the so-called *dilution effect*, the capacity of irrelevant information to dilute or weaken relevant information. Sometimes relevant information is called diagnostic, because it can help us make accurate predictions or diagnoses, and irrelevant information is said to be *nondiagnostic*. Using these terms, the dilution effect is the tendency for nondiagnostic information (like the description of Dick) to dilute diagnostic information (like the percentage of engineers vs. that of lawyers).

In this case the base rate of engineers is 70% and the base rate of lawyers is 30%. This information is highly relevant to the questions here. But descriptive information of marginal relevance can lead us to completely ignore highly relevant information about base rates.

Remember Mike (p. 335), the six two, muscular, aggressive college athlete? Why is it more likely that Mike is a banker than a pro football player? Because there are many more bankers than pro football players. The *base rate* for bankers is higher.

The **base rate** for a characteristic (like being a banker, or being killed by a pig) is the frequency or proportion of things in the general population which have that characteristic. It is sometimes called the initial or prior probability of that trait. For example, if one out of every twelve hundred people are bankers, the base rate for bankers is 1/1200. Often we don't know the exact base rate for something, but we still know that the base rate for one group is higher, or lower, than the base rate for another. We don't know the base rate for farmers or for chimney sweeps in the United States, but there are clearly far more of the former than the latter.

When we acquire information about someone or something (like our description of Mike) we need to *integrate* it with the old, prior information about base rates (many more people are bankers than pro football players). In the next section we will see that in many cases this can be done quite precisely. But the important point now is that although both pieces of information are important, in cases



where the size of the relevant group (or the difference in size between two relevant groups, e.g., bankers and pro football players) is large, the old, base-rate information can be much more important. Unfortunately, we often let the new information completely overshadow the prior information about base rates.

The **base-rate fallacy** occurs when we neglect base-rates in forming our judgments about the probabilities of things. We commit this fallacy if we judge it more likely that Mike is a pro football player than a banker (thus ignoring the fact that there are far more bankers than pro football players). Overreliance on the representativeness heuristic often leads us to underestimate the importance of base rate information. In the present case, Mike resembles our picture of the typical pro football player, so we forget what we know about base rates and conclude that he probably is one.

**Pigs vs. Sharks** We conclude this section with a quick examination of my student's claim that your chances of being killed by a pig are substantially greater than your chances of being killed by a shark. The claim is that live pigs, not infected pork that people eat, kill substantially more people than sharks do. Pigs vs. Sharks

The only way to be know for sure whether this is true is to check the statistics (if anyone keeps statistics on death by pig), but I would bet that my student is right. The base rate for contacts with pigs is *much* higher than the base rate for contacts with sharks. Most contacts are uneventful, but once in every several thousand, or hundred thousand, commonsense tells us that something will go wrong. So you probably are more likely to be killed by a pig, and it is much more likely that you will be injured by one. But a movie named *Snout* just wouldn't have the cachet of a movie named *Jaws*.

**Confusions about Inverse Probabilities** We know that a conditional probability like  $\Pr(\text{red} | \text{heart})$  may be quite different from its inverse, here  $\Pr(\text{heart} | \text{red})$ . The first probability is 1 whereas the second is  $1/2$ . But in many cases it is easy to confuse a probability and its inverse. It is true that the probability of someone fitting Mike's profile if they are a professional football player is reasonably high. By contrast, the probability of being a professional football player if they fit the profile is low (because the base rate of pro footballers is low, lower than the base rate of non-pros who fit the profile). Here it *is* easy to confuse a probability with its inverse. We will return to this problem in more detail in a later chapter.

### Safeguards

1. Don't be misled by highly detailed descriptions, profiles, or scenarios. The specificity makes them easier to imagine, but it also make them *less likely*.

2. Use base-rate information whenever possible. You often do not need any precise knowledge of base rates. Just knowing that there are a lot more of one sort of thing (e.g., bankers) than another (e.g., professional football players) is often enough.
3. Be careful to distinguish conditional probabilities from their inverses.

## 17.5 Anchoring and Adjustment

- Estimate the percentage of African countries in the United Nations (how much is it above, or below, 10%)?

The average response here is about 25% (the correct answer is 35%). But if you ask another group of people to

- Estimate the percentage of African countries in the United Nations (how much is it above, or below, 65%)?

the average response is about 45%. Why?

In the first case, most people think that 10% is too low, but they still *begin* with that figure and adjust up (to 25%) from it. In the second case people feel that 65% is too high, but they still begin with that figure and adjust downward (to 45%) from it. In each case their original starting point—10% or 65%—provides a reference point or *anchor*. We begin with this anchor and adjust up or down, but frequently we don't adjust enough. When we don't, the anchor has a strong effect on the judgment that they make. An anchoring and adjustment bias occurs when we don't adjust (up or down) enough from an original starting value or "anchor".

The anchor we use might be determined by the wording of questions or instructions, as it was above. But in different cases there will be different natural anchors that we'll tend to use.

Estimate, within 5 seconds, the product of:

$$8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$$

In experiments the median response is about 2250. But if you instead ask people to quickly estimate the product of:

$$1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$$

the median response is about 512. It appears that people perform just a few of the multiplications, *anchor* on the result, and *adjust* upward from there. In the first case the product of the first two or three digits is larger, so we adjust upward from a larger anchor and arrive at a larger number than we do in the second case. In this case neither anchor leads to a very accurate answer (the correct answer is 40,320).

We often fail to adjust enough to reference points or anchors

### 17.5.1 Anchoring Effects can be Very Strong

Anchoring effects can occur even when anchoring values are known to be entirely arbitrary, when they are ridiculously extreme, and when people are paid money for making correct estimates and predictions.

In the study that used our first example (involving the percentage of African nations in the United Nations), the anchor values were set by having each subject spin a wheel much like the one on *Wheel of Fortune* (it was rigged to stop at either 10% or 65%). So 10% or 65% served as anchors for the subjects even though they believed these numbers were completely arbitrary. But these anchors still had a strong impact on their estimates (in the 10% group the estimate was 25% and in the 65% group it was 45%).

Anchoring effects also occur when anchoring values are outlandishly high or low. When a group of psychologists asked subjects to estimate the number of Beatles records that made the Top Ten after first asking them if the number was less than 100,025, they found that this ridiculously high number served as an anchor and led subjects to give a high estimate (though not, of course, one anywhere near as high as the anchor itself). Even when people are offered money for doing well they remain susceptible to anchoring biases, and even the predictions of many expert forecasters will be influenced by arbitrary anchor values.

### 17.5.2 Anchoring and Adjustment in the Real World

In many situations the current situation—the way things presently are, the status quo—provides an anchor. In other cases first impressions provide an anchor. This may help explain the strength of the primacy effect. And some people think that anchoring helps explain hindsight bias (the tendency after the fact to think that we knew it all along). In hindsight we are anchored to what we know about how things actually turned out, and it's hard to think back and accurately reconstruct how we thought about things before we learned what the outcome was.

We can fall prey to an anchoring bias any time that examples or numbers are used to provide a frame of reference (“Estimate the number of people who live in Oklahoma; for example, is it between three million and four million?”). Often this happens without anyone intending to bias our judgments.

But whenever people are susceptible to a bias, there will be people who have learned to exploit that susceptibility. For example, experienced negotiators or people collecting for charities will often begin with extreme demands or requests in hopes of setting extreme anchors. Everyone knows that adjustments will be made in the direction of less extreme demands, but the more extreme the anchor, the greater its capacity to lead to an outcome closer to the one the negotiator wants.

Similar points apply when two people are haggling over a price. Other people in the persuasion professions, e.g., auctioneers and advertisers, can also exploit our susceptibility to anchoring effects by staking out extreme positions.

### 17.5.3 Safeguards

The most difficult thing is realizing that we are being influenced by an anchor at all. So the first step is to get in the habit of thinking about predictions and negotiations in terms of anchoring and adjustment. Once we do, we can look out for anchor values that seem too high or too low. We can also escape the power of anchors to color our thinking by considering several rather different anchor values. And if someone proposes an extreme anchor, counter with another anchor at the opposite extreme.<sup>1</sup>

## 17.6 Chapter Exercises

Answers to selected exercises will be found on page [345](#).

1. Debra is thirty two years old, outspoken, single, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and she has also participated in antinuclear demonstrations. Rank the probabilities of each of the following:

1. Debra is a bank teller.
2. Debra is either a bank teller *or* a leader in the feminist movement.
3. Debra is a bank teller *and* a leader in the feminist movement.
4. Debra is a bank teller, a leader in the feminist movement, and she plays the guitar.

2. Suppose that you don't know anything about Sue except that she is a woman. Rank the following from most probable to least probable:

1. Sue is a student
2. Sue is a student who wants someday to be the first astronaut from Rhode Island
3. Sue is either a student or a professor of English.
4. Sue is either a student or she isn't

---

<sup>1</sup>Additional references for some of the topics here will be found at the end of the next chapter. Further references to be supplied.

3. Defenders of O. J. Simpson correctly noted that only a very small proportion of husbands that beat their wives kill their wives. Why is this statistic likely to be accurate? What about the proportion of husbands who kill their wives among abusive husbands whose wives are later murdered? How do these two cases differ?

4. Couples often disagree, even argue, about who has been doing more than their share of the housework lately. Often each person quite honestly and sincerely thinks that they have done more than the other. This could happen without any self-deception or wishful thinking on the part of either person. How might some of the things we have learned in this chapter help explain it?

5. Recently I asked my class for examples of the way one of the heuristics we were studying could lead us astray. Someone mentioned that when people give us estimates for repairing our car or our home they might start out by giving us a high estimate. I realized that I had probably fallen victim to this just the week before, when a carpenter had given a very high estimate of the cost of a new roof. After hearing that I didn't feel so bad spending less than he recommended but a good deal more than I had originally planned. Which heuristic may have been involved here, and how did it work?

6. In the *Pretest* (p. 608) you were asked which alternative seems more likely in the next ten years:

1. An all-out nuclear war?
2. An all-out nuclear war that accidentally develops out a confrontation in the Middle East involving Iraq or Iran and some of their neighbors and that then spreads out of the region to other countries?

Which answer is right? Why? Which heuristic might aid and abet giving the wrong answer?

7. Do you think that there are more famous people from Oklahoma or from Kansas (a state of roughly the same size)? Don't proceed until you have thought about this question. How do you think people from Kansas would answer this question? What heuristics might be involved here?

8. Wilbur wants to buy a new car. He goes to the bookstore to get a mochaccino and the latest issue of *Consumer Reports*, which contains their most recent safety survey of full-sized sedans. The Volkswagen Passat has the highest safety rating, and Wilbur notes this. His next stop is to see his hairstylist, Wilma. Hearing that Wilbur is interested in buying a Passat, Wilma shrieks, "Those cars are dangerous! Total death traps. My client Suzy totaled hers when she ran into a utility pole going fifteen miles per hour!" So, Wilbur buys a Honda instead.

Why? What should Wilbur have done? What is the relevant heuristic or reasoning fallacy? How would you evaluate the credibility of his two sources of information?

9. We know the following about Wilma: She is 23 years old, athletic, and has taken various forms of dance for 18 years. Is she more likely to be a Dallas Cowboys cheerleader or a sales clerk for Dillards? Defend your answer and relate it explicitly to concepts and material studied in this chapter.

10. Which hand is more probable?

1. Ace of spades, king of spades, queen of spades, jack of spades, ten of spades.
2. Three of hearts, eight of diamonds, jack of spades, two of spades, nine of clubs.

Many people judge the first more likely. How might overreliance on one or more inferential heuristics lead to this error?

11. Wilbur is very shy and withdrawn, helpful, but with little interest in people or in the world of reality. He has a need for order and structure, and a passion for detail. Which is more likely and why: that Wilbur is a farmer or that he is a librarian? If you require more information to answer this question, what information do you need? How could you get it?

12. Suppose that we polled people and found the following (this example is hypothetical). The people polled were asked to estimate the percentage of American adults who were unemployed. Those who were employed underestimated the number, and those who were unemployed overestimated it.

13. In experiments subjects consistently err in judging the relative frequency of two kinds of English words. They estimate that the number of words beginning with a particular letter (for example 'R' or 'K') is greater than the number of words with those letters appearing third, even in the case of letters where words with the letter in third position are far more numerous than words in which the letter comes first. Explain what might lead subjects to this conclusion.

14. Most automobile accidents occur close to home. Why do you suppose this is true? Should you feel safer as you drive further from home?

15. Airlines sometimes post "full" fares prices that are higher than the fares they typically charge and automobile dealers often post suggested retail prices on the window sticker that are higher than they will actually charge. What all is going on

here? How effective do you think it is? What would be some good ways to resist things?

16. In December of 1989 Norman residents were warned about house fires. “On the average, we have two to three house fires with the Christmas season,” said Fire Marshal Larry Gardner, “and we haven’t had one yet.” Gardner appeared to be arguing that since we hadn’t had a severe house fire yet, we were very likely to have one soon. Assuming that this was his intention, what fallacy was he committing?

17. Earlier we learned about *illusory correlations*. Explain what an illusory correlation is and say how the availability heuristic might encourage us to believe in certain illusory correlations. Give an example to illustrate your points.

### Answers to Selected Exercises

10. The two hands are equally probable, but over reliance on the representativeness heuristic may lead us to think that the second hand is more likely because it is more similar to our mental picture of a random hand of cards.

12. In our example we imagined that a number of people were polled and asked to estimate the percentage of American adults who were unemployed. Those who were employed underestimated the number, and those who were unemployed overestimated it. Why? One reason is that unemployed people are more likely to live in areas and to go to places where there are other unemployed people, and of course their own situation is highly salient to them. By contrast, people who are employed tend to interact mostly with others who are employed too. Different samples were available to the people in the two groups.





# Chapter 18

## More Biases, Pitfalls, and Traps

**Overview:** In this chapter we study several more biases and tendencies to flawed thinking and consider ways to avoid them.

### Contents

---

<b>18.1 Framing Effects</b> . . . . .	<b>348</b>
18.1.1 Different Presentations of Alternatives . . . . .	348
18.1.2 Losses vs. Gains . . . . .	349
18.1.3 Loss Aversion . . . . .	351
18.1.4 The Certainty Effect . . . . .	352
<b>18.2 Psychological Accounting</b> . . . . .	<b>352</b>
<b>18.3 Magic Numbers</b> . . . . .	<b>353</b>
<b>18.4 Sunk Costs</b> . . . . .	<b>358</b>
<b>18.5 Confirmation Bias</b> . . . . .	<b>359</b>
<b>18.6 Self-Fulfilling Prophecies</b> . . . . .	<b>361</b>
<b>18.7 The Validity Effect and Mere Exposure</b> . . . . .	<b>362</b>
<b>18.8 The Just-World Hypothesis</b> . . . . .	<b>364</b>
<b>18.9 Effect Sizes</b> . . . . .	<b>364</b>
<b>18.10 The Contrast Effect</b> . . . . .	<b>366</b>
<b>18.11 How Good—or Bad—are We?</b> . . . . .	<b>368</b>
<b>18.12 Chapter Exercises</b> . . . . .	<b>368</b>

---

## 18.1 Framing Effects

### 18.1.1 Different Presentations of Alternatives

Which of the following two alternatives do you prefer?

- *Alternative A*: a 100% chance of losing \$50
- *Alternative B*: a 25% chance of losing \$200, and a 75% chance of losing nothing.

And which of A\* and B\* do you prefer?

- *Alternative A\**: an insurance policy with a \$50 premium that protects you against losing \$200
- *Alternative B\**: a 25% chance of losing \$200, and a 75% chance of losing nothing.

Most people prefer option B over A. And most prefer A\* over B\*.

But what is the objective difference between the two pairs of alternatives? The money comes out just the same with each pair. The only difference is that with the second pair of alternatives the loss is *described as insurance*. Whether we prefer risk or not is influenced by the way the risk and its alternative are described.

Before asking why this might be so, let's consider two more sets of alternatives. Imagine that the U.S. is preparing for the outbreak of a rare Asian disease that is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows.

- *Program C*: if C is adopted 200 people will be saved.
- *Program D*: if D is adopted there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved.

And which of the following two would you prefer?

- *Program C\**: if  $C^*$  is adopted, 400 people will die.
- *Program D\**: if  $D^*$  is adopted, there is a  $1/3$  probability that nobody will die, and a  $2/3$  probability that 600 people will die.

Most people prefer Program C over Program D. But most people (who have not seen the choices between C and D) prefer  $D^*$  to  $C^*$ .

Is this a problem? Yes, because the two pairs of alternatives are exactly the same; they are exactly equivalent in terms of how many people live and how many people die. As before, the only difference is in how the alternatives are described.

### Framing Effects

These differences in wording are said to **frame** the issue in different ways. When we frame a choice in terms of a certain loss we think about it differently than we would if we frame it in terms of insurance. When we frame a choice in terms of people being saved we think about it differently than we would if we frame it in terms of people dying. A small change in wording can have a big impact on our judgments.

**Framing effects** occur when the way we word or conceptualize alternatives influences which alternative people prefer. Such effects are often quite difficult to avoid; indeed, many people retain the choices they originally made in the above problems even after the contradiction is pointed out to them and even after they acknowledge it.

#### 18.1.2 Losses vs. Gains

In general, when probabilities of options are judged to be moderate to high, people are **risk averse** when it comes to potential gains. Being risk averse means that we avoid risky situations. We prefer a certain gain (say of \$10) to a 50/50 chance of getting \$20 (even though these alternatives have the same expected value). In fact, many people prefer a certain gain (say of \$10) to a 50/50 chance of getting \$25 or even more.

By contrast, when probabilities are judged to be moderate to high, people tend to be **risk seekers** when it comes to losses. Most of us prefer the risk of a large loss to a certain loss that is smaller. In our first example, most people prefer B (a 25% chance of losing \$200, and a 75% chance of losing nothing) to A (a 100% chance of losing \$50). How we think about risks often depends on how things are framed. More specifically, it depends on whether they are framed *as gains* (200 people are saved) or *as losses* (400 people die).

Whether we code events as losses or gains strongly influences how we think about them. Our examples so far have been artificial, but framing effects occur in the real world. For example, when several gas stations on the East Coast wanted to charge people more for using a credit card (because this entailed more expense for the gas station), their credit-card-using customers strongly objected to this “credit card surcharge”. The charge was framed as a penalty or a loss. But when the gas stations *reframed* the policy as a discount for using cash—which amounted to exactly the same thing in terms of the overall cost—their customers were more willing to accept it.

The study involving the fictitious Asian Disease provides an example of a *preference reversal*. It was theorized that it occurred because the first description of the two options presents or *frames* things in terms of a *gain* (saving lives) in relationship to the reference point (600 are expected to die), so that people are risk averse. But the second pair of descriptions frames the same options in a way that places the reference point at the status quo; here the options involve a *loss* (people dying), and so respondents are now willing to select what sounds like the more risky alternative. They reverse their preferences even though their options stay just the same.

In a more realistic study McNeil and his colleagues asked several hundred radiologists were given descriptions of two treatments, surgery and radiation therapy, for lung cancer. In half the cases the description was framed in terms of the cumulative probability of living longer than a given period of time. In the other half of the cases it was framed in terms of the probability of dying within that span (e.g., 85% chance of living longer than five years and 15% chance of dying within the next five years and ).

Surgery was preferred to radiation therapy 75% of time when it was put in terms of surviving, but only 58% of the time when it was put in terms of mortality (the major downside of surgery is dying soon afterwards and the dying frame may have emphasized that). Choices depended on whether the treatments were framed in terms of gains (people saved) or losses (deaths).

Lobbyists, trial lawyers, spin doctors, and public relations people are often quite skilled at framing issues in the way that is most favorable to their position. Unless we have strong feelings about the issue we often don't notice this, but when you are listening to such people it is always wise to ask yourself how the points they are making might be reframed.

### **Preference Reversals and Elicitation**

Some preference reversals (like that involving the Asian disease) involve framing effects. In other cases options are framed in the same way, but people's evaluations

are *elicited* in different ways. For example, if subjects are offered one of two options and asked to *choose* one of the pair, they tend to focus on the positive features of the two things. But when they are asked to *reject* one of the two (which leads to exactly the same result, namely getting one of the two things), they focus more on negative features (which is thought to be more compatible with the instruction to reject).

In a range of cases, responses seem tailored to be compatible with the statement of the problem or task, and this can lead to preference reversals. In such cases seemingly trivial differences in how we get people to express their preferences can get them to reverse their preferences. But few of us want our preferences to depend on trivial differences between ways of eliciting them.

### 18.1.3 Loss Aversion

Most people feel a particular aversion to loss. This *loss aversion* means that losses loom larger than corresponding gains. A loss of \$100 is more painful than the pleasure derived from a gain of \$100. Loss aversion at least partially explains two important phenomena: the *status quo bias* and the *endowment effect*.

The *status quo* bias is a bias in favor of the way things already are. Unless things are going badly, people often prefer to keep things the same rather than risk trying something new. Loss aversion helps explain this, since the potential disadvantages of changing things loom larger than the potential advantages.

We also value items we already possess (our "endowment") more than we would value them if we didn't have them. This is known as the *endowment effect*. We would typically require more money to sell something we already have than we would pay to buy it. This is clearly relevant to public policies involving regulatory takings (or eminent domain, where the government takes someone's land in order to build a highway, dam, or the like). Our aversion to loss explains the endowment effect as follows: once we have something, giving it up is seen as a loss, and so we require more in compensation for it than we would be willing to expend to acquire it.

The upshot of our discussion of framing and loss aversion is that the risks people are willing to take depend on whether they frame something as a potential gain or as a potential loss. Positive and negative frames lead us to think about things differently; different ways or framing the same options often influence peoples' preferences. It makes a difference whether alternatives are framed in terms of employment or unemployment; again, peoples' preferences are affected by whether options are framed in terms of crime rates or law-obedience rates.

Framing effects surely play a large role in politics and policy, where choices between uncertain options are ubiquitous and people with competing interests will

**Loss aversion:** a loss of a given size seems bigger than a gain of the same size

**Status quo bias:** wanting to keep things the way there are now

represent them in quite different ways. It certainly not completely true, but there is some truth in the old adage: it's not what you say—it's how you say it.

### 18.1.4 The Certainty Effect

**Certainty effect:** preference for certain outcomes over equally important uncertain outcomes

Would you pay more to reduce the probability of a serious disease from 90% to 85% or to reduce it from 5% to 0%? If you are like most people you would pay more for the latter. Although the objective decrease in risk is the same in each case, people have a strong preference for the “sure thing.” We prefer certain outcomes over uncertain ones.

In Russian roulette, for example, most people would pay more for a reduction of one bullet in the gun when it involves going from one bullet to none than when it involves going from two bullets to one (though if we were forced to play the game, most of us would spend a lot in either case). Given the option, most of us would choose a certain gain rather than take a chance on a larger gain that is only probable. For example, we would opt for a sure gain of \$1,000 over an 80% chance to win \$1,500 (or even more).

Equal probabilities are not always treated equally. If someone can frame an option in a way that seems to reduce all uncertainty about it, people will be more likely to accept it. For example, politicians who promise to completely solve a problem will fare better than those who merely offer policies that will probably make it less severe, even though the latter are frequently much more realistic. We don't like uncertainty.

## 18.2 Psychological Accounting

People frame the outcomes of choices as well as the choices themselves. Tversky and Kahneman asked people what they would do in the following situation:

**Case 1** You have decided to see a play where admission is \$10 per ticket. As you enter the theater you discover that you have lost a \$10 bill. Would you still pay \$10 for a ticket to the play?

88% of the respondents said they would still pay the \$10 for the ticket. Tversky and Kahneman then asked a number of other people about this situation:

**Case 2** You have decided to see a play and paid the admission price of \$10 per ticket. As you enter the theater you discover that you have lost the ticket. The seat was not marked and the ticket cannot be recovered. Would you pay \$10 for another ticket?

Only 46% would pay another \$10 for a ticket. But the two cases are completely equivalent in terms of the total amount a person would be out. Nevertheless, when the case is *framed* in these two different ways, people make quite different choices.

Tversky and Kahneman hypothesize that people do **psychological accounting**. In effect we keep different psychological books for different things. In Case 2 the respondents seem to add the second \$10 to the overall amount they would be spending on a ticket, and they aren't willing to pay \$20 to see the play. In this scenario people see the entire \$20 coming out of their "budget for play." But in the first case the \$10 they lost comes out of a different psychological account—not out of the account set aside for seeing plays—so they see the ticket as costing just \$10.

When we get windfalls, "easy money" we tend to think of it as less valuable than money we work hard for. An extreme example of this occurs when people are gambling. They tend to think of their winnings as "house money" which is not quite their own. So they find it easy to bet with it (and, typically, to lose it).

Many people owe thousands of dollars on their credit cards. The use of credit cards also involves something like mental accounting. It is easy to overspend on a credit card because it doesn't feel quite like spending. The money isn't coming out of a psychological account where we have stashed money we are reluctant to part with. In fact, however, money from this account is especially real, because we pay interest on our credit card bills. This is one thing people sometimes mean when they speak of 'compartmentalization'.

Although one of the best ways to save is to treat all money equally, we can sometimes use mental accounting to our advantage. For example, if you receive an unexpected gift, you can put the money in a bank account set aside to pay for next year's tuition. This moves it from the "easy-come easy-go" psychological account to another account that you are more reluctant to draw on.

## 18.3 Magic Numbers

We often judge outcomes and performances relative to some special number or target. Often such numbers were originally used as a way to measure something else, but sometimes they acquire a life of their own. They take on a significance that is disproportionately large compared to their actual value. Such variables are sometimes called *magic numbers*.

In the case of the economy the cost-of-living index or GNP may be treated as the measure of how well things are going economically. When evaluating a business with an eye to buying or selling stock price/earning ratio sometimes taken as the key. F. DeGeorge and co-workers found three variables of this sort in businesses: positive profit, previous year's earnings, and the agreement of financial analysts'

**Magic numbers:** numbers originally used to measure something else can take on a life of their own

earnings estimates. Closer to home, for many of you ACT or SAT scores and GPA easily become magic numbers for measuring academic promise and success.

Furthermore, we often set a somewhat arbitrary *threshold* for this number, with anything below the threshold counting as failure and anything above it counting as success. In public policy, business and other arenas where people want to measure progress or success, precise target numbers, quotas, acceptable level of risk, grade point average, and the like often come to define success and failure. The Local Charity will count it a failure unless they raise \$35,000, the new police chief fails unless she reduces crime by the targeted 2.5%.

Lawyers often speak of certain thresholds as *bright lines*. These are clear boundaries that make it easier to set policies people can understand. You have to be over a very clearly defined age to buy beer; a precise level of alcohol in a driver's bloodstream counts as driving under the influence. Bright lines are easy to judge (you aren't twenty one until midnight tomorrow), and they create clear expectations. We can generalize this notion to apply to threshold and boundaries in other settings; in the charity example raising \$35,000 is a bright target that tells us whether our fund-raising efforts were a success, or a failure

Target numbers are often ones that are salient and easy to remember; the local charity set a goal this year of raising \$35,000, not \$33,776. Furthermore, the numbers aren't completely arbitrary; no one would think raising \$3.50 was a worthwhile target. Still, within some vague range of sensible targets, the selection of a specific number typically *is* arbitrary.

Like many of the other things we have studied (e.g., heuristics), magic numbers and bright targets help us simplify very complex situations involving things like the economy and the environment. And frequently the numbers do tell us *something* about how the economy or a firm or a school is doing. Bright targets can also provide good motivation; if I set a goal of losing twenty pounds over the summer, it gives me something definite to shoot for.

In policy settings bright lines are sometimes useful because they provide a line that isn't renegotiated by each person on very occasion. In policy matters this is useful, because there is often a very real risk that the people making the decisions will do so in a way that isn't fair, or at least that won't appear fair. If the judge can simply decide if someone drank too much before getting behind the wheel, there would be a great deal of potential for abuse. Having a definite cutoff percentage of blood-alcohol level makes it more likely that everyone will be treated the same, and expectations are clear to everyone.

**Problems with Magic Numbers** Although magic numbers and bright targets are frequently useful, even unavoidable, they often take on a life of their own.



Originally they were a means to an end, e.g., an aid to seeing if the economy is improving, but eventually they become ends in themselves. This can lead to several problems.

1. Simply getting across the threshold is often seen as *the* measure of success, even when progress on either side of the line (getting closer to the goal by this much, exceeding it by that much) is equally important. Often the difference between almost making it to the target, on the one hand, and exceeding it by just a little, on the other, is insignificant. So a bright target can promote all-or-none thinking.
2. Target numbers can also be treated as the only relevant measures of success and failure, even though the precise target numbers are somewhat arbitrary and other, perhaps more important variables, are ignored.
3. In the worst case, magic numbers represent variables that are not very important, or their specific target values are not set in any sensible way.

**Magic Numbers and Suboptimal Policies** Magic numbers are often introduced in an honest effort to assess performance or progress, but they can have unintended consequences, including counterproductive policies and behavior. For example, policy makers sometimes establish bright lines when dealing with environmental issues.

For example, suppose that an agency sets a precise target for what counts as an acceptable level of arsenic in your community's drinking water. Clearly the arsenic level comes in degrees, with less arsenic being better, whatever the target value is. But a hard and definite number can foster a feeling that either risk is present or else it's eliminated.

It is often felt that anything short of a designated target is failure, and anything over it, even if it only limps over the line by just a bit, is a success. F. DeGeorge and co-workers found that people in charge of large businesses would often manipulate their earnings to get beyond a target value. Often it didn't matter how far from a target they ended up, as long as they passed it. Worse, it often didn't matter if the way of getting to the target would lead to problems in the longer run; for example, they would sometimes sell items at a large discount (or even a loss) late in the year, just to meet a target for yearly earnings.

C. Camerer and his colleagues found that New York City cabdrivers would set a target income for a day's work. They would drive until they reached that target, then knock off for the day. They would make more money with less driving if they drove fewer hours on days when business was slow and more hours on days when

it was brisk. But the daily-earnings target seemed to be a bright threshold that took on an intrinsic importance.

As a third example of how magic numbers can lead to less than optimal policies consider recent programs to make schools more accountable. The basic idea is a good one; schools have a very important obligation to their students and to those who pay the bills (in many cases taxpayers like us) to do a good job. Many recent efforts to judge how well schools are doing rely on standardized tests that are administered every few years.

In some cases this has led teachers to spend a great deal of time teaching the students things that will help them do well on such tests, while short changing other things. To the extent that the tests measure the all of the things students should be learning, this might be acceptable, but it is by no means clear that the tests do that. Indeed, we will see in the final chapter that if one goal is to foster critical reasoning, then training people to do well on standardized tests is not the best thing to focus on.

**Magic Numbers Can Become Ends in Themselves** Instead of being treated as indicators or indices of how well (or badly) things are going, magic numbers often displace what we originally cared about and become ends in themselves. We slip from thinking that GPA tells us something about how much a student is learning to thinking that GPA really *is* how well a student is doing. Furthermore such numbers can distort our picture of the situation. For example, one person may have a higher GPA than a second because he is taking much easier courses than she is taking.

There is also a political dimension to the selection of many magic numbers and targets. Someone in charge of a business or governmental agency will find it tempting to propose those measures of success that will show that they are in fact succeeding, while those who are unhappy with the way things are going may propose rather different magic variables to measure performance.

The allure of bright lines and magic numbers is one of several things that can lead us to focus on factors that are easy to measure or quantify, while paying less attention to things that are harder to measure, even when they are more important. This in turn can lead to the view that what can't be quantified or measured is unimportant, or even not real (we will return to this in Chapter 20).

Magic numbers and bright lines are often useful and even unavoidable, but they can lead to genuine problems. The key questions to ask when you hear some number cited as a sign of success or failure are (1) Is it really a good measure, an accurate indicator, of the thing that it's supposed to be measuring, and (2) does the specific target number associated with it really have any special significance?<sup>1</sup>

---

<sup>1</sup>James G. March's *A Primer on Decision Making: How Decisions Happen*, 1994, N.Y.: The Free

### Exercises

1. One of the middle schools in your city has been doing very badly by almost everyone's standards; students, parents, teachers, and neutral observers are all upset. Last week the school district allocated almost two million dollars to improve the school, and you are put in charge of a task force to assess how much (if any) the school really does improve. You will consult experts, but the experts you consult may disagree, and then your task force will have to come to a conclusion of its own. Think about the ways you would try to assess whether the tax dollars really translated into helping the students. Focus on some of the reasons why this would be difficult. How might people with different vested interests propose different measures of success? What might they be?
2. Management and workers at a large grocery chain are locked in a bitter negotiation over salary increases. What sorts of measures of how well the business is doing and how well the workers are doing might the workers focus on? What sorts of measures of how well the business is doing and how well the workers are doing might management focus on?
3. You are appointed as a student representative to a group asked to ascertain whether classes that feature a lot of group work provide a better educational experience than classes that don't. Your committee will consult experts, but the experts you consult may disagree, and then the committee will have to reach a conclusion of its own. How could you assess whether, on average, classes with a lot of group work did a *better* job of teaching people the things they should be learning than classes without group work do. Focus on the reasons why this would probably be difficult. How might people with different vested interests (e.g., those who spent a lot of time developing group projects vs. those who spent a lot of time developing lectures and individual projects) propose different measures of success? (We will return to some of the issues about group learning in Chapter 24.)
4. Give a real-life example of a bright line. Why is it supposed to be important? Who decided what it would be? How arbitrary do you think it is?

---

Press contains a brief but clear discussion of magic numbers. For a good discussion of bright lines in environmental policy see Anthony Patt and Richard J. Zeckhauser, "Behavioral Perceptions and Policies Toward the Environment," in *Judgments, Decisions, and Public Policy*, Rajeev Gowda and Jeffrey C. Fox, eds., 2001, Cambridge University Press. The quotation concerning bright lines and pesticides is from comments by Michael Hansen and Charles Benbrook for Consumers Union, July 17, 1997, on Proposed Reduced-Risk Initiative Guidelines, Docket Number OPP-00485. DeGeorge's work is presented in F. DeGeorge, J. Patel, and R. Zeckhauser, "Earnings management to Exceed Thresholds," *Journal of Business*, 1999, 72; 1-34. The study of cabdrives is C. Camerer, L. Babcock, G. Lowenstein and R. Thaler, "Labor Supply of New York City Cabdrivers: One Day at a Time," *The Quarterly Journal of Economics*, 1977, 112; 407-442.

5. What connections might there be between magic numbers, on the one hand, and issues involving framing and psychological accounting, on the other?
6. Think of some variable that is important but hard to quantify. How might it be overlooked if we focus too much on how it can be measured?
7. In arguing for the need for a bright line for what counts are a reduced risk pesticide, the Consumers Union and its Consumer Policy Institute (CPI), a non-profit product testing organization in Yonkers, New York, argued:

The purpose of these bright lines is to limit the pool of candidates for reduced risk status, and to provide EPA [Environmental Protection Agency] some easily applied criteria for saying "no" to requests for reduced risk status. Bright lines are also needed to reduce the time and agency resources required to evaluate requests for consideration for reduced risk status. If EPA opens itself up to a significant number of such requests, the time and resources entailed in reviewing and deciding upon such requests will end up diverting significant agency resources, and would hence defeat the purpose of the policy.

Restate the basic points in clear, non-technical English in a more general way that could apply to a number of issues in addition to pesticide labeling. How strong are such arguments for having bright lines in cases like this? What might be said in defense of the other side?

## 18.4 Sunk Costs

At the beginning of September you were really excited about OU's prospects in football, and when Wilbur suggested that you should each spend \$100 for a really good seat to the OU-Texas game in Dallas you enthusiastically agreed. But now it's October, the team hasn't done very well, and you aren't really very interested in football anymore. Besides, it's still very hot and you've had a cold and you'd really prefer to stay in Norman this weekend and relax. But then you remember that \$100. It's too late in the day to sell the ticket to someone else, so if you don't go to the game you'll have wasted it. Better get on the road.

We often reason this way, but is it rational? Your \$100 is *already gone*. Your overall financial situation is *just the same* as it would have been if someone had stolen that \$100 right before you bought the ticket. If *that* had happened, it wouldn't justify your going to the game. So how can the fact that you have sunk \$100 justify going now? It won't bring your original \$100 back. So how can it justify spending more money (for gas and food in Dallas) to do something you won't enjoy?

The \$100 Wilbur paid is known as a **sunk cost**. A sunk cost is money that has already gone down the drain. Since it is gone, it doesn't make sense to continue with a plan you no longer believe in simply because you sunk money in it. Following through on the plan won't bring that money back; it will only lead you to incur further costs (e.g., paying for gas and food on the trip to Dallas and having a miserable time doing it). Instead we feel like we must carry on in order to justify our initial expense.

**Sunk costs:** basing a decision on past investments (that are already gone) rather than on current prospects

Even people with a lot at stake often honor sunk costs. Wilma is the CEO of a large company that designs and manufactures attack helicopters. Her company has spent \$700,000 designing a new helicopter when it learns that a competing company has already designed a helicopter with the same features and has just landed a contract with the Pentagon. There won't be any other market for her company's helicopter. Should Wilma authorize another \$200,000 to finish the design project? If she does, she is honoring a sunk cost.

Sunk costs also operate at the national level. When a country is involved in a war that they aren't winning, one of the justifications typically offered to keep fighting, even when it can only lead to further disaster, is that "if we don't, all those soldiers who died will have died in vain."

Honoring sunk costs isn't always a bad thing to do. Sometimes it is important to us to follow through on a plan or commitment because we want to be the sort of person who finishes the things they start. And in other cases we can turn our tendency to honor sunk to our advantage. Many people reply on the power of sunk costs as a means of self-control, paying a good deal of money to join a health club or buy a home treadmill. Their hope is that the thought "all that money will go to waste if I don't go work out" will make them more likely to get to the gym.

## 18.5 Confirmation Bias

Many studies (as well as a bit of careful observation) document our tendency to look for, and remember, and acknowledge the value of positive evidence that supports our beliefs, while overlooking or undervaluing negative evidence that tells against them. The distinction between positive and negative evidence may be clearer if we consider a couple of examples.

**Confirmation bias:** tendency to look for confirming evidence while ignoring disconfirming evidence

For example suppose that you believe (or hypothesize) that

All swans are white.

Then a white swan is *positive evidence* for your view; it confirms or supports (though it doesn't conclusively establish) your belief. By contrast, swans that are not white are *negative evidence* against your view; they disconfirm your belief. In

fact, even one non-white swan shows that your belief is false. One black swan falsifies your hypothesis right then and there.

In many cases, though, negative evidence disconfirms a hypothesis without utterly refuting it. For example, Wilbur might wonder whether Wilma has a crush on him. The fact that she goes out of her way to chat with him is confirming, though by no means conclusive, evidence that she does. And the fact that she sometimes seems to avoid him is disconfirming evidence, though it doesn't prove that she doesn't.

*Confirmation bias* is our common tendency to look for, notice, and remember confirming or positive evidence (that supports what we think) while overlooking or downplaying disconfirming or negative evidence (which suggests that what we think is wrong). For example, if Wilbur is already convinced that women are bad drivers, he may be more likely to notice or remember cases where women drove badly and to overlook or forget cases where they drive well.

The *selective thinking* exhibited in this bias makes for bad reasoning, because it allows us to support our views without running the risk of finding out that they are wrong. It really boils down to considering only those things that support our own views, which is about as far from being open minded about things as we can be. Careful reasoning requires testing our views to see whether or not they fit the facts. The confirmation bias also encourages beliefs in illusory correlations (Ch. 15), since it encourages us to look for cases where two variables do go together without looking for cases where they may not.

Since the confirmation bias often leads to bad reasoning (that's why it's called a *bias*), it is important to avoid it. It should help to convince yourself of the value of negative evidence and to make a practice of looking for it, but this bias is difficult to eliminate.

In a series of experiments, Mynatt, Doherty and Tweney had subjects attempt to determine whether various laws or generalizations about the movement of a spot on a computer screen were true or not. They asked some people to *confirm* various generalizations about the dot's movement, others to *disconfirm* them, and still others just to *test* them. The instructions to disconfirm were not effective; about 70% of the time people in each group looked for confirming evidence. And the ineffectiveness of asking people to look for disconfirming evidence persisted even when the value and importance of searching for such evidence were explained to them before they began their task.

## 18.6 Self-Fulfilling Prophecies

A **self-fulfilling prophecy** is the tendency for a person's expectations about the future to influence that future in a way that makes the expectations come true. Sometimes we have expectations about things, most often about other people, that lead us, unwittingly, to treat them in a certain way. And treating them in this way may in fact lead them to behave in the way that we thought they would.

For example, if I hear that Wilbur is hostile before I ever meet him, I may be more likely to be hostile when I do meet him ("He's hostile, so I'd better beat him to the punch"). And this may lead *him* to react with hostility, even though he would have been friendly if I'd been friendly myself. My prediction leads me to act in a way that makes my prediction come true.

**Self-fulfilling prophecy:** tendency for a person's expectations about the future to influence the future in a way that makes the expectations come true

The psychologist Robert Rosenthal and his coworkers have studied self-fulfilling prophecies extensively. In a famous study in 1968, Rosenthal and Lenore Jacobson told grade school teachers at the beginning of the school year that their incoming students had just been given a battery of tests. Twenty percent of these students, it was explained, had great potential and should be expected to blossom academically in the coming year. In fact, the students in this group were selected randomly. Nevertheless, these twenty percent ended up improving more than the other students. What happened?

The chances of randomly picking the twenty percent that would improve are extremely small. Hence, the explanation is that teachers' expectations influenced their students' performances. Teachers expected the students in the targeted group to blossom, which led them to act in ways that encouraged the students to do so. For example, teachers gave the students in the high-potential group more time, more and better feedback, and more encouragement.

In short, the teachers' expectations led them to behave in ways that made their expectations come true. This sort of self-fulfilling prophecy is sometimes called the *Pygmalion effect*, after the play *Pygmalion* in which a professor of linguistics transforms a young woman with little education and bad grammar into a sophisticated, well-spoken person. Countless studies since have shown that this effect is very real (though often it is of modest size), both in the classroom and in other settings.

**Pygmalion effect:** people often perform better because we expect them to

Stereotypes can also serve as self-fulfilling prophecies. If teachers expect students from some groups to perform better than others, this may lead them to treat their students in ways that will make these expectations come true. In a society where people think that women are incapable of performing a demanding job like being a doctor, young girls are likely to be treated in a way that suggests they can't do such work. Furthermore, any interest they may display in medicine will be discouraged, and they will be encouraged to adopt quite different roles like being a



housewife. Years of such treatment *will* make it much more difficult for a woman to become a doctor. So the prediction that they can't be doctors can lead people to treat them in ways that will make the prediction come true.

## 18.7 The Validity Effect and Mere Exposure

### Validity Effect

**Validity effect:** mere repetition of a claim increases people's tendency to believe it

Researchers have found that the *mere repetition* of a claim will lead many of the people who hear it to think that it is more likely to be true (than they would have if they hadn't heard it before). This is called the **validity effect**: mere repetition makes the claim seem more likely to be true or more "valid." This effect occurs with true statements, false statements, and statements that involve expressions of attitudes.

In experiments on the validity effects subjects are often asked to rate the likelihood that a number of sentences (e.g., "Over 22% of the countries in the United Nations are in Africa") are true. In a later session they are asked to perform the same task, but with a partially overlapping set of sentences. On average sentences encountered in the first session receive higher rankings; subjects are more inclined to think they are true, simply because they have encountered them before. We seem to have a tendency to believe what we hear.

Since the validity effect can lead people to believe certain things without giving them any thought whatsoever, it is not surprising that it is exploited in propaganda, advertising, and related endeavors. If a company has enough money to run ads over and over, we will hear their claims about their product over and over. In many cases this will strengthen our tendency to believe those claims. The validity effect may also account for some of the biases and stereotypes people have. If you hear over and over how redheads are hot tempered, this will increase your tendency to believe it (especially if you don't interact much with redheads).

### Mere Exposure

**Mere exposure:** tendency to like things more simply because we've been exposed to them

There is an old saying that familiarity breeds contempt. The more you see of someone, the more flaws you notice, and you wind up thinking less of them. But in many cases this old saying is wrong. The more people are exposed to something (that they don't already dislike), the more they tend to like it. They don't need to interact with it, or hear it discussed. The **mere exposure** to the stimulus, without anything else happening at all, is enough to make them like it more.

In a standard experiment people (who don't know Chinese) were exposed to several Chinese hieroglyphics. Later they were shown a larger set of hieroglyph-



ics that included the ones they saw earlier as well as some new ones. The more previous exposures subjects had to one of the hieroglyphics, the more they liked it the second time around. Similar results have been obtained for many other sorts of stimuli. For example in one study subjects were first shown pictures of men's faces. The more times subjects saw a picture, the more they thought they would like the person.

Advertisers know about this. Not only can repeating their claims make them seem more valid (the validity effect). Simply exposing us to the name or a picture of their product can give us a comfortable sense of familiarity that translates into a purchase when we go to the store.

### Subliminal Mere Exposure

The mere exposure effect also occurs when people don't remember that they had previously encountered the stimulus. It even works, up to a point, when they were previously presented with a stimulus but weren't aware of it. Even when figures are flashed on a computer screen for a very brief period of time, too fast for subjects to be aware of them, the subjects will later show a preference for these figures over ones they haven't been exposed to previously.

In these cases the exposure is said to be *subliminal*; 'sub' means under and 'liminal' means consciously detectable, so something subliminal is something that can't be detected consciously (something supraliminal is something that can be). Are we susceptible to the subliminal influences of others? Can people manipulate our thoughts and actions by sending subtle, subliminal messages? In 1957 it was widely reported that a marketing group had conducted an experiment in a New Jersey movie theater. According to the story, messages like "Eat popcorn" and "Buy a coke" had been flashed on the screen, but so briefly that people in the audience weren't aware of them. And, the story continued, sales of popcorn rose by 58% and those of coke by over 15%.

No one wants to have their thoughts and actions manipulated by other people in this way, but fortunately there is no evidence that this story is true or that others can manipulate our thoughts and actions in such dramatic ways. We certainly are influenced by people's body language and tone of voice in ways we may not realize. But there is no evidence that advertisers or the manufacturers of "subliminal self-help" tapes are able to manipulate our thoughts and actions in any major ways.

## 18.8 The Just-World Hypothesis

We have a tendency to think that the world is fair and just. People usually get pretty much what they deserve and they deserve pretty much what they get. The psychologist Melvin Lerner called this phenomenon the **just-world hypothesis**: we think that things turn out, by and large, the way that they should. Life is basically fair. There is a good deal of evidence that many of us have a tendency to think this way. There are, of course, exceptions. Bad things (e.g., some terrible disease from out of the blue) do sometimes happen to good people. But when this occurs it often seems almost puzzling, unexpected. Typically, we tend to think, we reap pretty much what we sow.

**Just-world hypothesis:**  
tendency to think world is fair  
and that people get what  
they deserve

Lerner has shown that when people learn about an unfair outcome that is otherwise difficult to explain, they look for a way to blame the victim (“they must have done something to deserve this misfortune”). In an experiment by Ronnie Janoff-Bulman and her coworkers subjects heard a description of a young woman’s friendly behavior to a man she had met. They thought that her behavior was entirely normal and appropriate.

But other subjects who heard exactly the same description but were also told that she was then raped by the man thought that her behavior was more than friendly and that it encouraged the rape (“She was asking for it”). They *blamed the victim* (this is still not uncommon in cases of rape). Hindsight bias may also be involved here, since once an outcome (here the rape) is known, people often think they could have seen it coming.

We may want to believe in a just-world to make ourselves feel safer and more secure. As long as we do the right things, disaster probably won’t strike us (that wouldn’t be fair). But when we look at actual cases, we see that bad things can easily happen to good people and that people who aren’t so good can do pretty well. Good luck or bad luck can strongly affect things.

To the extent that we think this way, we will tend to think that most people who aren’t doing well are getting what they have coming. So if a group is treated badly, we may feel, they must have some defects that explain the bad treatment.

## 18.9 Effect Sizes

You hear on the news that a famous and trustworthy medical journal has recently published a study showing that a large daily dose of a moderately expensive dietary supplement, vitamin Q, will cut your risks of developing XYZ-syndrome *in half*. In other words, people who do not take the vitamin will be *twice* as likely to develop this syndrome. The syndrome is painful but not life threatening, and the vitamin

costs 65 cents a day. Should you start taking the vitamin?

The answer depends on a variety of things (e.g., whether you can afford the vitamin). But the first question you should always ask about such reports is:

What is the *base rate* of Syndrome XYZ?

Suppose that only 1% of the population (who don't take vitamin Q) ever develops XYZ-syndrome. Then if you take the vitamin, you cut your chances in half, down to 0.5%. In other words, you go from a 1 in 100 chance of developing the syndrome to a 1 in 200 chance. These numbers are so small that taking the vitamin may not be worth your time. By contrast, suppose that the base rate was 20%. If you could cut that in half, down to 10%, you would go from a 1 in 5 risk to a 1 in 10. The numbers here are big enough that you might want to give the vitamin some serious thought.

This is a fictitious example, but there are many real-life cases that illustrate the same point. Suppose that you learn that people who don't get enough vitamin C are ten times more likely to get botulism (which results from a deadly poison) or rabies. What are the first questions you should ask? What is the base rate for botulism? What is the base rate for rabies? It turns out that no more than three or four Americans die of either botulism or rabies in a given year. So even if some factor made rabies ten times as likely to kill you, your chances would still be about 30 out of 250 million.

But what if you learned of some precaution that could decrease your chances of having a heart attack by 20%? Again the relevant question is: what is the base rate? It turns out that about 1 in 3 Americans die from a heart attack, so if you could decrease your chances of heart disease by 20% it would be worth doing (we will return to this issue in more detail in the chapter on risks).

In these examples we are at least given percentages that tell us something about the impact of various drugs and the like. But media reports of experimental results often don't tell us about the magnitude of effects. The anchorwoman tells us that the manipulation of a certain experimental variable (e.g., taking vitamin Q) reduced cancer and that this result is *statistically significant*. But statistical significance does *not* mean the same thing as practical significance.

To say that a result is *statistically significant* simply means that it is unlikely that it was due to chance (to sampling variability). But with large samples, small and trivial differences are often statistically significant. For example a study might find that vitamin R reduces the risks for XYZ-syndrome by 0.20% (i.e., it reduces it by  $\frac{1}{5}$  of 1%). If our sample is large enough, this result may well be statistically significant. But the effect is so small that it won't be of much practical significance to anyone.

## 18.10 The Contrast Effect

Consider Figure 18.1. The two inside circles are exactly the same size, but the one on the right looks larger because of the size of the six circles surrounding it. In these two examples the *context* influences how we perceive things.



Figure 18.1: Circles in Context

Context can also influence how we *think* about things. The way that we think about or evaluate something often depends on the things around it. The alternatives, the points of comparisons, can strongly affect our perceptions, memories, judgments, inferences, and decisions.

**Contrast effect:** evaluations of, or judgments about, something are influenced by the contrast between it and things around it

The **contrast effect** occurs when our evaluations of, or judgments about, something are influenced by the contrast between it and the things around it. Many of our everyday judgments and inferences are affected by contrasts. George Bush would look short standing next to Shaquille O’Neal Humphrey (a tall basketball player), but tall standing next to a Shannon Miller (a short gymnast). The contrast effect is typically stronger the more similar the stimuli are to each other. For example, the effect is stronger when we compare Barry to two other people than when we have him stand beside two tractors of different heights.

When one thing is compared to something similar that is not as good as the first, the first thing is judged to be better than it would be without the comparison. In some cases both things are present at the same time, but the contrast effect also works when temporal contexts are involved. If the job applicant interviewed right before Wilbur does a terrible job, Wilbur is likely to seem better just by comparison.

A few years ago one of my colleagues taught two sections of the same course. The students in the first section simply refused to talk. Judging by the final grades, the second section wasn’t unusually good, but my friend soon came to think of them as very bright and talkative. The contrast with the first section made them look better.

We can exploit the contrast effect to make something look better (than it would have otherwise) by placing it in a context with something that looks worse. For example, a real-estate agent might show buyers an overpriced or dilapidated home

before showing them the home he wants them to buy. We can also make something look worse (than it would have otherwise) by placing it in a context with something that looks better. For example, the agent might discourage a person from buying a house by showing him a much better house first. And if you are in the market for a house, it is usually unwise to look at houses you know you can't afford. This will set up a contrast effect so that the houses you can afford won't look all that good.

### Other Context Effects

The wording of questions can affect our answers in many ways. Earlier we learned about a study where half the people in a group were asked "How frequently do you have headaches?" and the other half were asked "If you occasionally have headaches, how often?" The average response of the first group was 2.2 headaches a week while that of the second group was 0.7 headaches a week. Similarly, if you survey the people coming out of a movie and ask half of them "How long was the movie?" and the other half "How short was the movie?" those asked the first question will think the movie was longer.

The way options or possibilities are worded also influences people's responses to polls and public opinion surveys. For example, the results of polls to determine attitudes toward abortion vary depending on how the questions are worded. The questions in polls and surveys also often require you to select from a fairly restricted set of alternatives (e.g., should we increase defense spending or should we lower it?), which again tends to frame things in certain ways.

Surveys are often remarkably reliable, and if a number of surveys by different organizations converge on the same results, then we have good reason to believe them. But the wording effects we have encountered in this section should lead you to take any single survey with a grain of salt. This is especially true of surveys conducted by groups with a vested interest in the outcome. They can often make it more likely that they will find the response patterns they are looking for by framing their questions in ways that are likely to elicit the response that they want.

### The Compromise Effect

A good deal of research shows that many of us are reluctant to buy either the highest-, or the lowest-priced item. We prefer to "compromise" on a price somewhere in between. Businesses sometimes exploit this effect to sell more of one of their products.

For example, if Wilbur's factory has been selling two models of car stereos, one for \$200 and one for \$300, they may be able to increase the sales of the \$300 model by bringing out a \$400 model. The point isn't merely hypothetical. Researchers

have argued, for example, that Williams-Sonoma was able to sell more of their \$275 bread machines when they began producing a \$400 bread machine.

## 18.11 How Good—or Bad—are We?

In this chapter and the one before it we encountered a number of biases that can lead to bad reasoning. There is some debate about just *how* bad people are at reasoning. The heuristics-and-biases approach that figured prominently in the previous chapter was developed by Amos Tversky and Daniel Kahneman in a series of papers beginning in the early 1970s, and many people have found this approach very promising. But in the last few years some psychologists have argued that we are better at reasoning than some of this literature suggests.

Our performance depends, in part, on how we frame things. If we ask whether it is more probable that Linda is a bank teller or a bank teller and a leader in the feminist movement, we don't do very well. If we rephrase the question in terms of frequencies, rather than probabilities, we do better. If we ask: are there more people fitting Linda's profile who are bank tellers or who are bank tellers and leaders in the feminist movement, we give better answers. But while we may not be as bad at reasoning as some psychologists have suggested, it is clear that there is a lot of room for improvement.<sup>2</sup>

## 18.12 Chapter Exercises

1. Investors are often less willing to sell assets at a loss than they are to sell assets that have gained in value. Indeed many of us are very reluctant to see a falling stock. Is this sensible? What things should you consider in trying to decide whether to sell a stock or other asset that is losing in value.
2. People often go to investment counselors for help in investing their money. Should they put their money in stocks (with chances of a higher return, but more risky) or in bonds (less chance of high payoff, but less risk). So almost the first question an investment counselor asks a new client is how much risk they can live with. How might subtle differences in the way the counselor words this question affect the answers she'll receive (and, hence, the advice that she will give)?
3. Although you can't arrange such things, do you think you might do better in a

---

<sup>2</sup>For a discussion of the validity effect see H. R. Aries, et. al. "The Generality of the Relation between Familiarity and Judged Validity," *Journal of Behavioral Decision Making*, 2 (1989); 81-94 Further references to be supplied.

job interview if the person right before you turned in a terrible performance? Why? How would you test a hypothesis about this?

4. Most of us have heard that taking an aspirin a day decreases our chances of getting a heart attack. This result is statistically significant. But is it of any practical significance? What would it mean to say that it is? What information would you need to answer this question? How would you get it? Get it.

4. What role might the endowment effect play in insurance fraud? Explain. Given an example or two, and be sure to defend your answer.

5. The following was a question on the pretest: Which do you prefer?

- a. a 100% chance of losing \$50
- b. a 25% chance of losing \$200 and a 75

About 25% of you chose (a), and 75% of you chose (b). Explain why the majority of you probably picked (b). Be sure to use to relevant concepts that were discussed in class. What type of attitude regarding loss does this reveal about the majority of us?

6. Countless Americans are currently battling a credit crisis. Indeed, many people owe thousands of dollars in credit card bills. Explain what psychological accounting is and how it could contribute to this crisis. What remedies does psychological accounting suggest?

7. Give an example, preferably one from your own experience, of the compromise effect. How susceptible to it do you think most of us are?

8. Researchers randomly divided a group of Duke undergraduates into two groups. The first group was asked to imagine that they had a ticket to that years NCAA Final Four tournament (a more important item at Duke than many other universities). They were asked how much it money it would take them to sell their ticket, and the median price was about \$1500. The second group was asked how much they would pay for a ticket to the Final Four and their median price was about \$150. What could account for this enormous difference? Do you think these figures would hold up if first group of students actually had the ticket? Can you think of similar examples? What do they tell us about human reasoning?

8. You and your friend both paid \$8.00 to see the movie, but by the time it is half over you both realize that you aren't enjoying it at all, and it only seems likely to get worse. What reasons are there to stay? What reasons are there to leave? What



should you do? What would you do? What do you say when Wilbur asks how you can waste your hard-earned money by leaving early.

9. I often find that once I've installed a piece of software on my computer's hard drive I'm very reluctant to remove it, even if I never use it. After all, I tell myself, I might need it one of these days. But when I recently reinstalled the operating system I found it fairly easy not to put some of the pieces of software back on my computer. What might be going on here? How sensible is it? Can you think of similar examples in your own experience?

10. You have lost weight on your new diet but a few minutes ago you broke down and got a Big Mac and large order of fries at the McDonalds' drive-through window. You haven't eaten them yet. What should you do?

11. "To terminate a project in which \$1.1 billion has been invested represents an unconscionable mishandling of taxpayers' dollars." – Senator Jeremiah Denton, 11/4/81. Is this true or not? What more would you need to decide. What problems could be lurking here?

### Answers to Selected Exercises

1. Refusal to sell falling stock or other assets like a house that is declining in value is often a futile effort to honor sunk costs. You would be better off deciding which stocks to buy and sell based sensible expectation of their future performance, not on what has happened in the past. If there is good reason to think the stocks or the housing market will rebound, hold on to them. If not, cut your losses.

8. This is a case of sunk costs.



Figure 18.2: Spin



## Chapter 19

# Cognitive Dissonance: Psychological Inconsistency

**Overview:** Earlier we studied the notion of *logical inconsistency*. In this chapter we will study *psychological inconsistency*, peoples' *perceptions* of their own inconsistency, and the ways these influence their attitudes, beliefs, and thoughts. We will see that although making our beliefs and attitudes more consistent is typically a good thing to do, one strategy for doing this (dissonance reduction) often results in bad reasoning.

### Contents

---

<b>19.1 Two Striking Examples</b> . . . . .	<b>372</b>
<b>19.2 Cognitive Dissonance</b> . . . . .	<b>373</b>
19.2.1 How Dissonance Theory Explains the Experiments . . .	374
<b>19.3 Insufficient-Justification and Induced-Compliance</b> . . . . .	<b>375</b>
19.3.1 Induced-Compliance and Counter-Attitudinal Behavior	375
19.3.2 Prohibition . . . . .	376
<b>19.4 Effort Justification and Dissonance</b> . . . . .	<b>377</b>
<b>19.5 Post-Decisional Dissonance</b> . . . . .	<b>377</b>
<b>19.6 Belief Disconfirmation and Dissonance</b> . . . . .	<b>378</b>
19.6.1 When Prophecy Fails . . . . .	379
<b>19.7 Dissonance Reduction and Bad Reasoning</b> . . . . .	<b>379</b>
<b>19.8 Chapter Exercises</b> . . . . .	<b>380</b>

---

## 19.1 Two Striking Examples



Last week you signed up to be a subject in a psychology experiment. Now you walk into the psychology lab, sit down across the table from the experimenter, then notice a large plate of fried grasshoppers in front of you. After some initial discussion of other matters, the experimenter asks you to *eat* a few of the grasshoppers. What would you do?

The pressure to comply with experimenters in situations like this is much greater than is often supposed, and many of the subjects in this 1965 study by Philip Zimbardo and his coworkers actually ate several grasshoppers. But the experimenters manipulated what turned out to be a very interesting variable; they randomly assigned each of the subjects to one of two groups.

**Nice-experimenter Group:** In this condition a warm, friendly experimenter nicely asked subjects to eat grasshoppers as a favor.

**Cold-experimenter Group:** In this condition a cold, aloof experimenter pressured subjects to eat grasshoppers.

The subjects were later asked (by a third person) how much they liked the grasshoppers. No one was wild about them, but which group do you think disliked them the least? It turned out that the group that had been asked by the aloof experimenter had a more positive attitude toward eating the grasshoppers than the group that had been asked by the friendly experimenter.

A 1959 study by Festinger and Carlsmith sheds some light on this puzzling outcome. They asked each of their subjects to perform a boring, repetitive, meaningless series of manual tasks—arranging and rearranging rings on spools—for an hour. They then asked each subject to go outside to the waiting room to tell the next subject how interesting and enjoyable the experiment was and to remain on call to talk to other subjects about it, in the event that the experimenter's assistant would be unable to do so. In other words, they asked the subjects to *lie*.

The subjects were randomly assigned to two conditions:

**High-Reward Group:** Subjects in this condition were paid \$20 to lie to the person waiting outside.

**Low-Reward Group:** Subjects in this condition were paid \$1 to lie to the person waiting outside.

Subjects were later asked how much they had enjoyed the hour-long task. Now that you know the outcome of the grasshopper experiment, you may be able to predict what they said. The high-reward (\$20) group felt that the activity was very dull. It

was dull, so no surprise there. But the low-reward (\$1) group felt that the task had been more interesting. What's going on?

In each case, the group that had a *strong external inducement* to do something they didn't want to do (eat grasshoppers, lie to the waiting subjects) didn't change their original attitude (about eating grasshoppers or about how boring the task was). But the group that had a *weak external inducement* did change their attitude (the grasshoppers weren't *really* so bad; the task really wasn't all *that* boring).

## 19.2 Cognitive Dissonance

Festinger devised Cognitive Dissonance Theory to explain such phenomena. This theory was very popular in the 1950s and 60s, was studied less in the next two decades, but has been making a comeback in the 1990s. Festinger argued that when a person perceives *inconsistencies* among her actions, attitudes, and beliefs she will experience an unpleasant motivational state that he called 'cognitive dissonance' ('cognitive' means 'psychological' and 'dissonance' means 'disharmony', so the idea is that the person feels a disharmony or conflict among their beliefs, attitudes, and the like). Dissonance is psychologically uncomfortable.

**Cognitive dissonance:** a state of tension when one sees their actions, beliefs or attitudes as inconsistent

The notion of dissonance will be clearer if we contrast it with two other notions. Some of my actions and attitudes *reinforce* one another: I oppose gun control, and I belong to the NRA; I support campaign finance reform, and I voted for the candidate who supports it. Others are *irrelevant* to one another: I oppose gun control, and I brush my teeth. But some of my actions and attitudes are *psychologically inconsistent*: I believe smoking can kill me, but I smoke two packs a day; I think lying is wrong, but I lied through my teeth to get this job. Such inconsistency will often produce cognitive dissonance.

**Cognitive dissonance** is an emotionally unpleasant state of tension that results from such perceived inconsistencies. For example, telling a lie to the waiting subjects (action) seems inconsistent with my view that I'm not the sort of person who would tell a lie unless there was a really good reason to do so (belief).

Cognitive dissonance involves tension and discomfort, so people will try to *eliminate*, or at least *reduce* it. The way to reduce it is to typically to modify some of ones actions, beliefs, or attitudes. Since past actions have already occurred and a person cannot change what has already been done, dissonance reduction will typically involve a change in attitude or beliefs. This will be easier to see if we consider how dissonance theory explains the two experiments described above.

### 19.2.1 How Dissonance Theory Explains the Experiments

In both experiments subjects are induced to do something they don't want to do. Eating grasshoppers is disgusting and lying to the person outside is wrong. In order to explain such phenomena, dissonance theory requires one additional assumption:

When we have strong external reasons or justification for doing something that we don't approve of, we can explain why we did that thing by noting this justification.

Subjects in the high-reward condition of the second (boring task) experiment could reason this way (though they didn't do so consciously): I told a lie. I think lying is wrong and I'm not the sort of person who lies unless there is a good reason to. But *sometimes* there *are* good reasons. For example, it is acceptable to tell a little white lie to avoid hurting someone's feelings ("How do you like my new haircut?"). That wouldn't really show that I'm deceitful. Similarly, in this case, I had a good external reason to tell a lie (the \$20). In short, subjects in this condition could conclude that the lie didn't really reflect badly on them, because they had a strong external justification (\$20) to tell it.

But subjects in the low-reward condition didn't have this out. They could only reason this way: I told a lie. I think lying is wrong and I'm not the sort of person who would tell a lie unless there was a good reason to do so. But I didn't have a good reason (\$1 isn't enough to justify it). So these subjects feel an inconsistency among their beliefs and actions: I lied; I wouldn't lie without a good reason; I didn't have a good reason. The result: cognitive dissonance.

Festinger reasoned that subjects who lied for \$1 couldn't really justify doing it for so little money. So in order to avoid seeing themselves as deceitful—to make their action *consistent* with their attitudes—they (subconsciously) modified their attitude toward the experiment. It really wasn't as boring as they originally thought.

The pattern of explanation of the second experiment is the same. Subjects who encountered the friendly experimenter had a good external justification for eating the grasshoppers. They were doing something in order to help a nice person that they liked. But subjects who had the unfriendly experimenter couldn't justify their actions in this way. They were stuck with some dissonant views about themselves: I just ate those disgusting grasshoppers; I don't do things like that without a good reason; I had no good reason to eat them. To reduce this inconsistency, they modified their attitude. The grasshoppers weren't really *that* disgusting after all.

In this chapter we will examine four types of situations where cognitive dissonance plays a role in our actions and thought. The first, which we have focused on thus far, involves induced-compliance.

## 19.3 Insufficient-Justification and Induced-Compliance

### 19.3.1 Induced-Compliance and Counter-Attitudinal Behavior

These experiments illustrate the first of three types of insufficient justification effects that we will consider in this chapter. It is sometimes known as the **insufficient-justification through induced-compliance paradigm** of dissonance reduction.

In both studies subjects were *induced* to do things that they didn't really want to do. The experimenter got them to engage in "counter-attitudinal" behavior (i.e., to do things that ran counter to their attitudes—like telling a lie). But in each study half of the subjects were induced to do so with what seemed to them like very weak justification. These subjects could *not* find a good external justification for doing what they did, and this produced cognitive dissonance between the counter-attitudinal behavior and the attitude itself. Since the subject could not go back in time and undo the behavior, the only way to reduce this dissonance was to modify their attitudes so that they become more consistent with telling the people outside that the experiment was interesting (the task wasn't really *that* boring).

Such shifts in attitude are known as *insufficient-justification effects* because they arise when justification or coercion is so small that it seems to the subject insufficient to justify her behavior. Thus, subjects seemed to find \$1 an insufficient justification to lie and to find eating grasshoppers at the behest of someone they didn't like an insufficient justification for eating them.

Note that the justification is *in fact* sufficient to get the subject to do something they don't want to do (since most of them did eat grasshoppers or lie). But later it seemed so mild that it was difficult for the subject to realize that this was what had led them to do what they did. They saw the justification as insufficient. They were subtly pushed to do something, but it felt to them like they freely chose to do it.

Not all examples of attitude change in response to induced compliance are trivial. For example, European-American students were asked to write essays in support of large scholarships for minority students (which many of them opposed). Half of the subjects were told that the exercise was voluntary (low external incentive). The other half were told that it was required (high external incentive). The subjects with the high external incentive didn't change their attitudes about affirmative action, but the subjects with a low external incentive developed more positive attitudes to minority students. Similar results have been found for attitudes toward many other topics, including police brutality and the legalization of marijuana.

When people experience cognitive dissonance they will typically try to modify the inconsistent element that is least resistant to change. So although one of Zimbardo's subjects could theoretically reduce dissonance by denying that she ate the grasshoppers, it is obvious that she just did, and so it is easier to change her

When we feel we lack sufficient justification for doing something that runs counter to our attitudes we may modify our attitudes

attitude about eating grasshoppers. Again, deeply held views that enhance one's self-esteem will be more resistant to change than many of our more peripheral, less deeply-held attitudes.

Insufficient-justification effects leading to attitude change have been found in a very wide range of conditions. They are especially strong when the following conditions are met (but there is good evidence that dissonance and attempts to reduce it can arise even when they are not met).

1. The person sees the counter-attitudinal behavior as freely chosen (if it was coerced, then the coercion would explain the behavior).
2. The behavior could be foreseen to have some bad consequence.
3. The person sees himself as responsible for these consequences.

### An Alternative Explanation: Self-perception Theory

If I'm doing X, I must not think X is so bad

Daryl Bem proposed an alternative account of such phenomena. He argued that people discover their own attitudes and emotions partly by observing how they actually behave. When internal cues are ambiguous or hard to interpret, we are in much the same position as an outside observer who is trying to interpret us. According to Bem, subjects in the two experiments inferred their attitudes by observing how they actually behaved. Thus, subjects paid \$20 inferred that they lied because they were well paid. But subjects paid \$1 inferred that they said what they did because they believed it (since there being no strong external reasons to justify it).

It remains a matter of controversy whether Bem's account or dissonance theory's account provides a better explanation of these two experiments (there is some evidence that Bem's theory is right about certain types of cases and dissonance theory is right about others).

We won't worry about this issue here, however, since the phenomena themselves are what matter for our study of reasoning. We will speak of these sorts of results as *dissonance results*, and we will see that in many cases it is quite plausible to suppose that peoples' aversion to perceived inconsistency plays an important role in their thought and behavior.

### 19.3.2 Prohibition

A related type of insufficient-justification involves prohibition. In a 1963 study Aronson and Carlsmith told nursery-school children that they could not play with an attractive toy. Half the children were threatened with a mild punishment if they played with the toy; the other half were threatened with a more severe punishment.

Later the children in the mild-threat condition valued the toy less than the children in the severe-threat condition.

Dissonance theory's explanation is that the children in the severe-threat condition had a very good external justification not to play with the toy. They could have said to themselves: I like the toy, but I don't want to be punished and that is why I'm not playing with it. But the children in the low-threat condition couldn't reason this way. The threat was very mild, and so it provided *insufficient* justification for avoiding the toy. This led to dissonance: I like the toy; I play with toys that I like; But I'm not playing with this one. They reduced this dissonance by devaluing the toy. ("It's really not that attractive after all"; what implications might this have for getting children—or adults—to change their attitudes?)

## 19.4 Effort Justification and Dissonance

The second sort of dissonance phenomena involves our need to justify the effort that we put into something. For example, numerous studies show that when people undergo a severe or difficult initiation to join a group, they value membership in the group more than people who don't have to go through so much to get in.

People who undergo a severe initiation seem to reason as follows: I am a sensible person who would not put myself through this difficult initiation if it were not worth doing; I am putting myself through all this difficulty. Now if the group isn't worth belonging to, this package of thoughts is inconsistent and that will create dissonance. But if the group is really worth the effort, that could justify what I've been going through. More generally (just as your grandmother always said), people tend to value things more when they have to work hard to get them.

## 19.5 Post-Decisional Dissonance

The third sort of dissonance phenomena involves decision making. We often have to make difficult choices between alternatives: Where should I go to college? What should I major in? Which job offer should I accept? Should I marry Wanda? Should we put off having children until we are more settled?

In a difficult decision each alternative has some pluses and some minuses, and we aren't sure how to balance them out in a way that will lead to the best choice. You are trying to decide whether to bring the collie or the terrier home from the Norman Animal Shelter. Both dogs have pluses and minuses. The collie seems smarter, but she may be too big for your little apartment; the terrier is cute, but seems a little dumb and you've heard terriers are difficult to house train.



Whichever dog you choose, you will give up some positive features (of the dog you don't take) and accept some negative features (of the dog that you do take). Your awareness of these positive and negative features will be dissonant with the choice that you actually made. This is known as **post-decision dissonance**: after a difficult choice we are likely to experience dissonance.

Post-decisional dissonance is greater when the choice is hard to undo, because we can't reduce the dissonance by changing our decision. In such cases how could we reduce it? Once people commit themselves to a choice, they often exaggerate both the positive aspects of the thing they chose and the negative aspects of the thing they rejected. Once you chose the terrier, you may conclude that a collie would have been too much trouble, probably wouldn't have been affectionate, and that terriers are much smarter than you had supposed.

This strategy for reducing post-decisional dissonance shows up in many studies. Jack Brehm posed as representative of a company that was doing consumer research on household products. He asked people rate the desirability of various household appliances like coffee makers and toasters. As a reward for participating in the study, each woman was offered a choice between two of the items that she had rated.

Later the women were asked to re-rate the desirability of the products. Brehm found that the appliance the woman had chosen was rated higher than it originally had been, while the appliance she could have chosen, but didn't, was ranked much lower. This is known as the *spreading effect*; we often feel like there is a greater difference between the desirability of things after we choose between them than we did beforehand.

Although the evidence is less clear cut, some of it suggests that after a difficult decision people also often become selective in the information they seek about the things they chose between. They seek out and attend to information that supports their decision (after bringing home your terrier, Wilbur, you read about the virtues of terriers) and avoid or discount information that doesn't support it (you quit reading about the strong points of collies).

## 19.6 Belief Disconfirmation and Dissonance

The fourth dissonance phenomenon, and the last one we will study, involves the disconfirmation of someone's belief (a belief is *disconfirmed* when there is clear evidence that it is false). Information that is inconsistent with our beliefs can produce dissonance. This can lead us to avoid the information, or to ignore it, or to dismiss, or to attack the people who convey it to us (we have encountered all of these strategies before).



Sometimes, however, it becomes so obvious that the belief is false that tactics like these simply will not work. The disconfirmation of a belief can produce dissonance, since we felt like it was true, and if we took action on the basis of the non-disconfirmed belief, the possibilities for dissonance are especially strong. We will now consider a very interesting, real-life example of this.

### 19.6.1 When Prophecy Fails

In 1954 Leon Festinger came across a newspaper account of a small “doomsday” cult who believed that the world would end on December 21. His coworkers infiltrated the group and observed the members’ behavior. The group members were very committed to their beliefs. They had gotten rid of all their possessions (who needs a toaster when the world is about to end?) and were genuinely preparing for the world to end.

December 21 came and went, and the world didn’t end. This dramatically disconfirmed the group leader’s prediction, and we might expect that the members of the group would have lost their faith and left. Members of the group who were alone on December 21 did lose their faith, but those who were with the rest of the group did just the opposite. They concluded that their own actions had postponed the end—though it would arrive soon—and this seemed to *strengthen* their faith.

Before their belief had been disconfirmed, members of the group hadn’t done much to convince others to join them, but after the disconfirmation, they worked hard to convert others to their own position. Their new belief, that their actions had delayed the end of the world, restored the consistency between their belief in the head of the group and the fact that they had given away everything they owned, on the one hand, and the fact that her prophecy failed, on the other.

## 19.7 Dissonance Reduction and Bad Reasoning

Over the short run, dissonance reduction often allows us to see our views as well-founded and our actions as compatible with our ideals, but it doesn’t make for good and independent reasoning. Indeed, some of the ways dissonance reduction works will be familiar to students of fallacies. When someone offers arguments for views we don’t like or evidence suggesting that we are wrong or that our actions are harmful, there are several common ways of reacting. All of these can help us reduce dissonance.

1. Distort the person’s position or argument or evidence so that we don’t have to take it seriously (e.g., the strawman and either/or fallacies).

2. Shift the focus away from the person's position or argument or evidence so that we don't have to think about it (e.g., the ad hominem, red herring fallacy).
3. Overestimate the quality of the arguments or evidence supporting one's own position (e.g., appeal to a suspect authority, appeal to ignorance).
4. Rationalize that "everybody does it," so we might as well too.

While such strategies may protect our attitudes and self-image, being unwilling to confront the facts does not promote clear and independent thinking. In the following exercises we will encounter a number of examples of the importance of cognitive dissonance and dissonance reduction in the real world.<sup>1</sup>

## 19.8 Chapter Exercises

What role do you think that cognitive dissonance and attempts to reduce it play in the following cases (answers to selected exercises are found on page 381)?

1. Many Jews in Germany and other European countries saw signs of terrible danger as the Nazis came to power. But many of them made little effort to leave.
2. Many of the people who worked in the concentration camps saw themselves as good, decent human beings, even after the war was over. How could this be?
3. Suppose you were a heavy smoker when the Surgeon General's report about the dangers of cigarettes came out in 1964. How would you react? List several ways that a smoker might try to reduce dissonance when learning about the report.
4. Recall the different perceptions students had of the Princeton-Dartmouth football game. How might dissonance reduction be involved in the very different interpretations they had of this game?
5. The more difficult it is to become a member of a group (e.g., because it costs a lot of money, because of harsh hazing practices), the more people who do become members tend to value it. Give an example of this. How does dissonance enter the picture?
6. Suppose you are strongly tempted to cheat on the final for this course. Once you have decided what to do, you will probably experience some dissonance. Why? How might you reduce it?

---

<sup>1</sup>For a very readable discussion of cognitive dissonance theory by one of the people who played a major role in its development see Elliot Aronson's *The Social Animal*, N.Y., 1992. The book, *When Prophecy Fails*, by Leon Festinger, H. W. Riecken, and Stanley Schacter, University of Minnesota Press, 1956, is easy and fascinating reading. Further references to be supplied.

7. If people change their attitudes more when they do things for small rewards, what might effects might punishment have on attitude change?
8. Suppose that some person or group has already invested a lot in something (money keeping a car going, lives lost in a war). There is some tendency to think that justifies further investment. Could dissonance theory be relevant here?
9. In his excellent discussion of dissonance theory, Elliot Aronson says that a modern Machiavelli might well advise a ruler:
  1. If you want someone to form more positive attitudes toward an object, get him to commit himself to own that object.
  2. If you want someone to soften his moral attitude toward some misdeed, tempt him so that he performs that deed. Conversely, if you want someone to harden his moral attitudes toward a misdeed, tempt him—but not enough to induce him to commit the deed.

What do you think about this advice?

10. Wilbur is struggling to decide between buying a house and renewing the lease on his apartment. There are positive and negative factors on each side. If he buys the house, he will have a tax deduction on his mortgage and he will be building up equity in something that he owns. But he will have to care for the lawn, and he is financially responsible for things that break. On the other hand, if he renews his apartment lease, someone else cares for the lawn and fixes things when they break. But he won't be getting a tax write-off or building up any equity. After considerable agonizing, Wilbur decides to buy the house. How is Wilbur likely to reason, and feel, after he makes his decision?
11. In §19.6 we examined the role that cognitive dissonance might play in leading us to ignore disconfirming evidence or to attack those who present it. First explain how this works. Then explain how the reduction of such dissonance is related to the strawman fallacy, the ad hominem fallacy, and confirmation bias.

### Answers to Selected Exercises

1. Many Jews in Germany and other European countries saw signs of terrible danger as the Nazis came to power. But many of them made little effort to leave.

It would be very hard to get out of the country, and it is difficult to reconcile your worst fears with many other beliefs you have. And once you have decided to stay, post-decision dissonance reduction is likely.

2. Many of the people who worked in the concentration camps saw themselves as good, decent human beings, even after the war was over. How could this be?

One common finding is that those who worked in the camps came to see their victims as less than human. When you treat someone badly, there is a tendency to derogate them, to think “well, they deserved it.” How could this reduce dissonance?

3. Suppose you were a heavy smoker when the Surgeon General’s report about the dangers of cigarettes came out in 1964. How would you react?

The report was careful and thorough, a good authority. But you might be inclined to disbelieve it (otherwise you would face the dissonant thoughts: I care about my health; smoking is bad for me; I smoke.) A study done at the time showed that only 10% of nonsmokers doubted the report. 40% of heavy smokers did.

4. Recall the different perceptions students had of the Princeton-Dartmouth football game. How might dissonance reduction be involved in the very different interpretations they had of this game?

Our beliefs and values influence what we focus on and how we interpret it. Seeing things in ways that fit with our views is one way to reduce (or prevent) dissonance.

5. The more difficult it is to become a member of a group (e.g., because it costs a lot of money, because of harsh hazing practices), the more people who do become members tend to value it.

“I went through hell to get into this group. It must be worth belonging to.”

6. Suppose you are strongly tempted to cheat on the final for this course. Once you have decided what to do, you will probably experience some dissonance. Why? How might you reduce it?

Judson Mills did a study of cheating with sixth graders that helps answer this question.

1. Those students who had succumbed to the temptation developed a more lenient attitude toward cheating.

2. Those students who had resisted the temptation developed a more negative view about cheating. (Why?)
7. If people change their attitudes more when they do things for small rewards, what might effects might punishment have on attitude change?

Aronson and his coworkers found that mild threats of punishment were more effective in changing attitudes than harsh threats.

1. They hypothesize that if a person does something solely because they fear a severe punishment, they don't come to change their attitudes about it. They do it *because* the punishment would be severe, not because of the attitudes that they happened to have.
2. If they do something when there is only a mild threat of punishment, they can't so easily explain their behavior by the present of a strong external inducement..



## **Part VIII**

# **Evaluating Hypotheses and Assessing Risks**





## **Part VIII. Evaluating Hypotheses and Assessing Risks**

In Chapter 20 we examine several key notions: evidence, prediction, testing, causation, experimentation, and explanation. We consider the ways these three notions function in science, in pseudoscience, and in our everyday lives. Parts of this chapter are under construction.

In Chapter 21 we study risks, the misperception of risks, and ways to more accurately assess the riskiness of various actions and projects.



## Chapter 20

# Causation, Prediction, Testing, and Explaining

**Overview:** Science is a complicated human practice, and there are many different sciences. No simple account can do justice to all aspects of every scientific field, but we will examine the main features that are present in most. We will also see that many of the tasks that confront the scientist also confront ordinary people when they attempt to understand the world around them, including other people and even themselves. In such situations, ordinary people can be thought of a “intuitive scientists.” We conclude with an examination of pseudoscience and superstition.

The section on intuitive vs. statistical prediction and the appendix are complete; the rest of this chapter is under construction

### Contents

---

<b>20.1 Science</b> . . . . .	<b>390</b>
20.1.1 Getting Data . . . . .	391
20.1.2 Testing Hypotheses and Predicting the Future . . . . .	391
20.1.3 Tracking Down Causes . . . . .	393
20.1.4 Mill’s Methods . . . . .	396
<b>20.2 Experiments</b> . . . . .	<b>400</b>
20.2.1 Controlled Experiments . . . . .	400
20.2.2 Were the Results due to Chance? . . . . .	402
<b>20.3 Giving Explanations</b> . . . . .	<b>403</b>
<b>20.4 The Everyday Person as Intuitive Scientist</b> . . . . .	<b>405</b>
20.4.1 Gathering Data . . . . .	406

20.4.2 Testing and Predicting . . . . .	407
20.4.3 Tracking Down Causes . . . . .	409
20.4.4 Giving Explanations . . . . .	409
<b>20.5 Intuitive vs. Statistical Prediction . . . . .</b>	<b>409</b>
<b>20.6 Pseudoscience . . . . .</b>	<b>414</b>
<b>20.7 Chapter Exercises . . . . .</b>	<b>417</b>
<b>20.8 Appendix: Scientific Notation and Exponential Growth . . . . .</b>	<b>419</b>
20.8.1 Scientific Notation . . . . .	419
20.8.2 Exponential Growth . . . . .	422

## 20.1 Science

The key feature of genuine science is that its claims are subject to test, in one way or another, but various other things also play a central role. No simple account can do justice to every aspect of every scientific field, but we will examine several features that are central to most of them. These include:

1. Getting data (sampling)
2. Drawing Inferences from sample to population
3. Assessing covariation (correlation)
4. Formulating Theories or Hypotheses (often ones about what causes what)
5. Testing those Theories
6. Prediction
7. Explanation

Some items on this list are more important in some sciences than in others, but much science involves some version of most of them.

We have studied the first three topics in earlier chapters. The fourth topic, formulation of hypotheses, often involves creativity and there is no recipe for devising good hypotheses anymore than there is a recipe that would always enable you to write good songs, though we will say a bit about it below. We then turn to the last three tasks in more detail.

The list is not meant as a flowchart of what a scientist does. It would be hopeless to begin from scratch and simply collect a vast amount of data. Science is not done in a vacuum, and almost all scientific work begins with problems and questions that are or important for practical reasons (how can we cure an HIV infection?) or for theoretical reasons that depend on the current state of the science

(how much dark matter is there in the universe?). But along the way the scientist will be involved with the various activities on the list.

We will use the word ‘hypothesis’ to mean the same thing as ‘theory’. Sometimes when we call something a ‘hypothesis’ or ‘theory’ we mean to suggest that it is dubious, or even false. But as we will use these two words there is no suggestion of this sort. A hypothesis or theory is a claim that we are often concerned to test.

### 20.1.1 Getting Data

#### Drawing Inferences from Samples to Populations

We studied samples and populations in Chapter 15, and we can quickly recall the basic points here. We often infer a conclusion about a population from a description of a sample that was drawn from it. When we do

1. Our *premises* are claims about the *sample*.
2. Our *conclusion* is a claim about the *population*.

For example, we might draw a conclusion about the percentage of people who favor sending troops to Kosovo from premises describing the responses of 700 people to a poll on the subject. In such a case our inference is *not* deductively valid. It involves an inductive leap.

But if we are careful in our polling, our inference can still be inductively strong. This means that if we begin with true premises (which in this case means a correct description of the sample), we are likely to arrive at a true conclusion (about the entire population).

A good inductive inference from a sample to a population requires:

1. A large enough sample
2. A representative (unbiased) sample

How do we apply the things we learned to scientific inference? . . . . .

### 20.1.2 Testing Hypotheses and Predicting the Future

Outline of the basic concepts that will be discussed here.

1. A theory or hypothesis is *testable* just in case some sort of objective, empirical test could provide evidence that it is true or else evidence that it is false. A theory or hypothesis is *falsifiable* just in case some sort of objective, empirical test could show that it is false.

2. One tests a hypothesis in science, medicine, auto mechanics, detective work, etc. by making a *prediction* and then seeing whether that prediction is true or not. Theories can (often) predict and explain outcomes and phenomena, but they cannot be deduced from data or outcomes.
3. A prediction typically takes the form: if such and such *test conditions* are realized, then such and so should result.
4. In an *experimental test*, we can bring about the test condition in a laboratory or in some other controlled setting. Experimental sciences employ many experimental tests, but in some sciences (e.g., astronomy, meteorology) such tests are difficult to devise.
5. An *auxiliary hypothesis* is an background assumption used in testing a theory or hypothesis of interest. Every test of any interesting scientific theory involves auxiliary hypotheses (e.g., about the workings of the measuring devices one employs, the presence or absence of various disturbing influences, and so on).
6. Not all tests are equally good tests of a theory. In general a difficult or *severe test* of a theory is much better than a weak or easy test. The more unlikely a prediction seems to be before we actually check it, the better the test it provides of a theory.

For example, if your local weatherperson has a meteorological theory that predicts it will rain in Seattle sometime this coming April, we won't be bowled over if this comes true (we all knew that it would rain at least a little sometime during April, long before we ever heard of her theory).

But suppose her theory predicts that it will rain between nine and nine and a half inches in Seattle between noon and 1 AM on April 7. If this happens, we are surprised and take it to provide strong—though not conclusive—support for her theory: it must have *something* going for it, to get something like this right. Other things being equal predictions that are extremely definite and precise provide a better test of a theory than predictions that are indefinite or vague.

7. *Variety of Evidence*: In general having many different kinds of evidence is much better than having a lot of evidence of the same kind. The more unlikely a prediction seems to be before we actually check it, the better the test it provides of a theory.
8. An *ad hoc hypothesis* is a hypothesis that one makes up after a test disconfirms a theory one likes (“after the fact”), solely to save the pet theory. It has no independent support.
9. A *crucial experiment* is an experiment that pits one theory against another.
10. *Theories can't be conclusively proven* because the general form of a test is

not one of a valid argument.

11. *Theories can't be conclusively falsified* because when predictions don't turn out correct, the result can be pinned on one of the auxiliary hypotheses.
12. But the two previous points do not mean that theories cannot be tested. They can, and we often have very good reasons not accept one as approximately correct or to reject another as just plain false.

### Assessing Covariation (Correlation)

We discussed correlation in Chapter 15. The points to add here are the following:  
.....

### Formulating Hypotheses

...

1. Context of Discovery
2. Context of Justification

....

### 20.1.3 Tracking Down Causes

.....

### Type-Causation vs. Token Causation .....

### “The” Cause vs. Causal Factors .....

### Causation and Correlation

Correlations often point to causes; indeed, they are often vital *evidence* for claims about what causes what. When two variables, like drinking and liver disease, tend to covary we suspect that there must be some reason for their correlation—surely *something* must cause them to go together. But correlation is not the same thing as causation. For one thing, correlation is symmetrical (smoking and heart attacks are correlated with each other), but causation is a one-way street (smoking causes heart attacks, but heart attacks rarely cause people to smoke). So just finding a positive correlation doesn't tell us what causes what.

But as we saw in Chapter 15, correlation is not the same thing as causation. Correlation  $\neq$  Causation  
For one thing, correlation is symmetrical (smoking and heart attacks are correlated

with each other), but causation is a one-way street (drinking causes liver disease, but liver disease rarely causes people to drink). So just finding a positive correlation doesn't tell us what causes what. But it's a start. What are the next steps?

### Bad Causal Reasoning

As with other sorts of reasoning, causal reasoning can go awry in many different ways, but there are several patterns of defective causal reasoning that are common enough that we should discuss them here.

**Post hoc, ergo propter hoc** The Latin phrase still turns up often enough that it's worthy learning it. It means: after this, therefore because of this. We commit this fallacy when we conclude that *A* caused *B* because *B* followed *A*.

When we put it this way it is likely to seem like such hopeless reasoning that it isn't really worth warning about it. Day follows night, but few of us think that day causes night. There are many cases, however, where it really is tempting to reason in this way. For example, we sometimes take some action, discover the outcome, and conclude that our action led to the outcome. We encountered numerous cases of this sort when we learned about regression to the mean (p. 317).


For example, if the institution of a new policy is followed by a decrease in something undesirable or an increase in something desirable, it may be tempting to conclude that the measure caused for the shift. The crime rate went up last year, we added more cops or passed tougher sentencing laws, and this year it came back down to its normal level. In many cases this return to normal might have occurred without the measure, simply as a consequence of regression to the mean. In such cases, we are likely to *explain* the drop in crime by the increased number of police or the new laws, but we will be wrong and the new measure will be given credit it doesn't deserve.

To take another example, unusually good performances are likely to be followed by less outstanding performance simply because of regression, and unusually bad performances by better ones. If we neglect the possibility of regression effects, this may lead us to suppose that criticizing someone for a bad performance is a more effective way to getting them to do well than praising them for a good performances is.

As a final example, when some people recover from a given illness if they take a certain drug it is tempting to conclude that the drug caused their recovery. But might they have gotten better anyway, without it? In many cases people do. Here we need to compare the rate of recovery among those who take the drug with the rate among those who do not. In many cases, as we will see in a bit, the best way to do this is with a controlled experiment. The connection here may just be an



illusory correlation (p. 295), and if it is, there is no interesting causal connection here at all.



**Common Causes** When two sorts of things are positively correlated it is sometimes the case that one causes the other. But we reason badly when we take a correlation between two things to show that one causes the other. In the chapter on samples and correlations we noted that correlations between two things are often based on some third, common cause. For example, there is a positive correlation between a falling barometer and a rain storm, but neither causes the other. They are the joint effects of a common cause: an approaching cold front.

**Separating Cause from Effect** Sometimes the problem here is described as confusing cause and effect. How could anyone have a problem with this? Isn't it usually obvious. Yes, very often it is. But in complex systems it is often difficult to determine what causes what.

Families with a member who is schizophrenic tend to be dysfunctional in various ways. But does the schizophrenia lead to familial problems or do the problems lead to the schizophrenia? The answer needn't be simple. In addition to these two possibilities, it may be that there is some third, common, cause. Or it may be that each thing makes the other worse. There may be a sort of vicious circle with a feedback loop.

**Feedback Loops** .....

**Confounding Variables** In other cases a number of things go together in ways that may make it difficult to determine exactly what causes what. For example Afro-Americans receive worse health care, as a group, than whites. But it is sometimes argued that this isn't a direct result of racial discrimination but of economic differences (which often are a result of discrimination), lack of good health insurance, or yet other factors.

.....

## Causal Schemas .....

## 20.1.4 Mill's Methods

**Mill's methods:** five strategies for tracking down causes

Mill's methods are techniques designed to help us isolate the genuine causes from a list of potential causes. They are often called *eliminative* because they pinpoint the true cause (when they do) by eliminating potential causes that aren't genuine. These techniques get their name from John Stuart Mill (1806–1874), the British philosopher who systematized them. They weren't invented by Mill, though, since people who engaged in careful causal reasoning always employed similar strategies.

Mill's methods are not magical. They cannot pinpoint causes in a vacuum. They require us to make substantive assumptions about sorts of things that might have caused a given event or type of event. In many situations we can do this with a reasonable degree of confidence, so the methods are often useful. But they are not foolproof; causal reasoning is inductive reasoning, and so it is subject to standard inductive uncertainty.

We will present a modern version of Mill's methods here, rather than the one he set out in 1843. We will approach things in terms of necessary and sufficient conditions, so you might want to quickly review the discussion of them in §3.2.

**Method of Difference**

**Basic Idea** In an effect of interest occurs in one case but not in a second, look for a potential cause that is present in the first case but not in the second. If you find one, it is the cause (or at least part of the cause).

**Simple Examples** Getting your car to start; debugging a computer program; making a better cake; allergic reaction.

**Details of the Method** We focus on cases that are different. The effect of interest, e.g., my car's starting, occurs in one case but not in the second.

**More Complex Applications****Simple Illustrative Diagram**

$$A B C \Rightarrow E$$

$$\neg A B C \Rightarrow \neg E$$

.....

### Method of Agreement

**Basic Idea** If a potential cause  $A$  is present when an effect  $E$  is absent, then  $A$  isn't the cause of  $E$ . If we can eliminate all but one of the potential causes in cases where  $E$  occurs, then the remaining potential cause is the actual or genuine cause.

#### Simple Examples

**Details of the Method** Here we focus on cases that are alike, that agree.

#### More Complex Applications

##### Simple Illustrative Diagram

$A B C \Rightarrow E$

$A D C \Rightarrow E$

$A B D \Rightarrow E$

.....

### Joint Method of Agreement and Difference

The joint or combined method simply put the first two methods together.

**Basic Idea** If  $A$  is always present when, but only when,  $E$  is present, then  $A$  caused  $E$ .

**Simple Example** All of the people who took the Wilburton Prep course for the LSAT (a test used in making decisions about admissions to law schools) were accepted by the law school they wanted to attend. But none of the people who applied to the same law school but didn't take the course were accepted.

#### Details of the Method

#### More Complex Applications

##### Simple Illustrative Diagram

.....

### Method of Residues

#### Basic Idea

#### Simple Examples

**Details of the Method****More Complex Applications**

.....

**Method of Concomitant Variation**

This method is related to our earlier discussion of correlation in § 15.3. It is relevant when different amounts or rates of something are involved. For example, the higher a person's blood pressure the more likely they are to suffer a stroke.




Figure 20.1: Concomitant Variation

**Basic Idea**

**Simple Examples** The more cigarettes I smoke the more headaches I have.

**Details of the Method****More Complex Applications**

The word 'concomitant' is not as common as it was in Mill's day, but it simply means accompanying. This method applies when an increase in one variable is accompanied by an increase in the other.

Figure 20.1 on the facing page depicts hypothetical data indicating that the higher a person's average caloric intake per day the higher their chances of being overweight. Note that it is a simple and natural extension of the type of diagrams we used to represent correlations.

.....

### Exercises on Mill's Methods

Explain which versions of Mill's methods are involved in the following cases and assess the plausibility of the arguments that make use of them.

1. Wilbur: In recent years several states have enacted right to work laws (laws that make it illegal for workers to have to join a union). In all of these cases states were soon collecting more tax dollars. Our state budget is about to be cut and tuition fees at state schools are about to go up yet again. So we need to pass a right to work law here.

2. Wilma: That's not right. In some states where a right to work law was passed state tax revenues increased, but in other states that passed such laws tax revenues actually fell. We do need a better budget, but this shows that passing a right to work law isn't the answer.

.....

In the following exercises list some potential causes of the effect that is singled out and explain how you might use one or more of Mill's methods to try to pinpoint the actual cause.

3. You have successfully raised tomatoes each of the last four years, but this year almost all of your crop is going bad, with tomatoes dying before it's time to pick them.

4. Your four year old twins, Wilma and Wilbur go to the same day care center and much of the rest of their time together. But Wilma has come down with measles while Wilbur has not. Why? How might Wilma have gotten it? How did Wilbur escape.

.....

### Case Study: Semmelweis and the Cause of Childbed Fever

Tracking down causes is sometimes very difficult: often much work goes into finding the cause of the victim's death, the origin of a disease, or who committed a crime. Here is one real-life example.

The Hungarian physician Ignaz Semmelweis worked at the Vienna General Hospital from 1844 to 1848. During this period babies were delivered in two wards or divisions of the hospital. In the First Division a distressingly large percentage of mothers, it averaged about 8.5% a year, died of childbed fever. The death rate in the Second Maternity Division was substantially lower at about 2.3% a year. Indeed, even mothers who didn't get to the hospital and gave birth at home or elsewhere had a much better chance of not getting the disease.

There is a serious practical question here: what caused so many more deaths' in the First Division? The answer wasn't obvious, and it took a good deal of detective work by Semmelweis to discover it. Here is what he did.

.....

## 20.2 Experiments

### 20.2.1 Controlled Experiments

The basic idea behind simple experiment is straightforward. We divide subjects into two groups, treat one group in one way, the other in a different way, and see if the differences in treatment *causes* a difference of the sort we are interested in. For example, we might give one group the new drug, give the other group no drug, and see if those who get the drug start feeling better. This is a sophisticated use of the method of difference.

Things are more complicated than this basic idea might suggest, however, because of sampling variability (p. 282), i.e., because each sample we draw from a population is likely to be at least a little different from most other samples.

Here we will take a more detailed look at the simplest sort of experimental set-up. More complex experimental designs are possible, but most of them simply involve more elaborate versions of the key ideas we will discuss here. In the simple case we can think of the subjects in the experiment being divided into two groups. It is important that they are *randomly assigned* to their group. What this means is that some random process, e.g., use of a random number table, is used to place people in groups. Indeed, we want more than this; we want each group of subjects to be just as likely to be assigned to the experimental group as to the control group.

The point of all this is to avoid biasing our results. For example, if we assigned all the people who came in during the morning to one group and those who came in during the afternoon to another, our two groups might differ in important ways that we could avoid.

Random assignment is *not* the same thing as random sampling. The participants in the subject are seldom a random sample of any group of much interest. It is the people who will show up to be in the study, perhaps college Freshmen who




Figure 20.2: Basic Experimental Design

agree to be subjects to earn extra credit in Psychology 101.

The key notions here are *control group*, *experimental group*, *independent variable(s)*, *dependent variable(s)*, and the difference between *experimental results due to chance*, on the one hand, and *results due to manipulation of the independent variable*, on the other.

### What is the Population?

It is easy to identify the sample in an experiment (or a field study, for that matter); it's just the people (or things) in the study, those who answered our questions, or the like. But we don't do experiments to just to find out about the things in our sample. We want to use what we learn about a sample to make an inductive inference about a population. In some cases it is clear what the population is: in a poll, for example, it might be the class of all potential voters. But in many cases it is not.

What, for example, is the relevant population if we run an experiment using subjects conscripted from Psychology 101 (which is what happens in many psychology experiments)? Is the relevant population, to which we want to generalize, all of the people enrolled in psychology classes, or all those at the university where the study was done, or all people of roughly the same age as the subjects, or those with roughly the same income, or . . . . We only need ask such questions, which could go on almost forever, to see that there is no one, simple answer.

The problem here is related to the problem of the reference class (p. 576). The more broadly we construe the population (e.g., all American adults), the more interesting our findings will be, since they apply to a very large group. But this is

risky, because we have sampled such a select group (students in Psych 101) from this huge population. It is safer to generalize to smaller populations, but this is much less interesting (“My study shows that Freshmen enrolled in Psych 101 here where I teach . . .”). There is tradeoff between how safe and how interesting our claims are. One danger here is to being with the wider, more interesting claims, but to retreat to the more narrow ones if someone presses us on the matter.

In some case, for example, medical research, it matters a lot what the relevant population is. The people in our study did well when given a certain drug, but it is safe to generalize to people who seem to be rather different? Is it safe to treat them with the same drug? Common sense and knowledge about the relevant field (e.g., medications for heart disease) may help us here.

In the end we could really only get a precise picture of the population that our results generalize to by seeing how widely how experiment replicates. This means running the experiment with a number of different groups and seeing whether we obtain the same result with them. If we do, they are part of the population. But this is a long and costly process, and one point of doing an experiment is to avoid all that time and expense. So in the end we are often uncertain about the population, though the more we know about the field in question, the better our hypotheses about it are likely to be.

## 20.2.2 Were the Results due to the Manipulation or to Chance?

### Sampling Distributions and Null Hypotheses

This will be less technical in next draft

What makes experimentation really difficult, especially with subjects that differ greatly among themselves (as people do), is *sampling variability*. If we drew many samples from the same population, we are likely to get somewhat different results each time. This is why experimenters have to rely on statistics. The tools they use can get a bit complicated, but here we will only aim for a very general view of the structure of an experiment. We will continue to focus on experiments with just two groups, but most of the points here also apply to more complex experimental designs.

Statistical significance  $\neq$  practical significance

.....  
Null hypothesis significance testing is confusing for nonspecialists. Perhaps the most important thing about it, for nonspecialists, is that *statistical significance* is a mathematical concept and it need not mean substantive or practical significance. Tiny and trivial effects may, especially with very large numbers of subjects, be statistical significant. So when you read that the results in a given study were significant, it may not mean that they were large or important.

There are several things that make it difficult to love null-hypothesis signif-



ificance testing (the problems have been laid out clearly many times now; e.g., ..... example, you know going in that a non-direction null hypothesis is in fact almost certainly false (any two populations will differ a little) and (for similar reasons) directional nulls will be false half the time. But the chief problem is simply this. Even if you get a  $p$ -value that allows you to reject the null hypothesis, this does not tell you what you really want to know.

Let  $E$  stand for the fact that you got an effect at least as large as the one that you got. In practice an “effect this large” is measured by a  $t$  statistic (or  $F$ -ratio of the like), but this in turn is a function of effect size (e.g., mean differences) together with sample size, so here we can simply speak here of an effect at least as large as the one you got. So you have run your subjects and analyzed your data and now know  $E$ . You want:

1.  $\Pr(H_o | E)$ , how likely  $H_o$  is given your result?

Instead you get:

2.  $\Pr(E | H_o)$ : how likely you are to get  $E$  if  $H_o$  is in fact true.

These are very different probability (compare: the probability that you draw a red card given that you draw a king is  $\frac{1}{2}$ , but the probability that you drew a king given that you drew a red card is only  $\frac{2}{25}$ . You could get 1. from 2. by Bayes' theorem. But you would also need to know  $\Pr(H_o)$  and  $\Pr(E)$  (or their equivalents), and you don't know these; in the latter case especially it's not all that clear what the unconditionalized  $\Pr(E)$  even means.

So lacking these, 2. does not tell you the probability that the null hypothesis is true or false, or the probability that the alternative is true or false. It also doesn't tell you the probability that your result is replicable. Nor does it say much about the size, much less any real significance, of the effect; it only tells you that, given your sample size, the effect was big enough to get you out to the critical region where you could reject the null hypothesis.

.....

## 20.3 Giving Explanations

We are constantly trying to make sense of things: we need to *explain* and *understand* the world around us. Almost every time we ask *why* something happened or *how* something works we are seeking an explanation. This is true in almost every part of science and in every aspect of our daily lives.

One of the main goals of science is to explain the world around us. For example, disease has always been a serious affliction of human beings. In ancient times

diseases were often attributed to supernatural causes, e.g., demons, but such theories did not provide very effective ways to treat or prevent disease.<sup>1</sup> The work of Louis Pasteur and others toward the end of the nineteenth-century led to the germ theory of disease. This theory allows us to understand the causes of many diseases and to explain why they spread in the ways they do. And this understanding in turn led to vaccines and other measures that allowed us to eliminate some diseases and curtail the spread of others.

Learning about things and understanding how they work is often rewarding in and of itself, and it is vital if we are to deal successfully with the world around us. If we understand how things work, we will be able to make more accurate predictions about their behavior, and this will make it easier for us to influence how things will turn out. If you understand how your computer works, you will be in a much better position to fix it the next time it breaks down.

### **The Explanation Reflex: Telling More than we can Know**

We are constantly seeking explanations in our daily lives. My car started just fine yesterday; everything seems the same today, so what explains the fact that it won't turn over now? We are particularly concerned to understand the behavior of other people. *Why* did Patty Hearst join the group that had kidnapped her and treated her brutally (p. 477)? Why did several hundred people at Jonestown commit mass suicide? But sometimes our desire to make sense of things is so strong that it leads us to see patterns or causes in places where there aren't any.

This happens in the case of illusory correlation, for example, where we think we see relationships even when there aren't any. We saw a very dramatic example of this in the Appendix to Chapter 8, p. 154). Subjects there gave reasons for their behavior that clearly couldn't be right, but they were completely convinced that they were.

Other studies point to similar conclusions. Nisbett and Wilson found that people would concoct reasons (subconsciously and after the fact) for choosing between two things that were exactly alike. In another study by Langer and her coworkers, people were much more likely to respond to a request if the person making it gave them a reason for doing it (You should do it *because*, ...), even when the reason was completely irrelevant to the request.

If someone simply asked to step in front of a person making copies at a Xerox machine the person was much more likely to comply if they were given a reason for the request. In some cases the reason was a good one ("Because I only need a few copies and I'm in a rush). But people complied almost as often when they

---

<sup>1</sup>The abbreviation 'e.g.' means 'for example'; 'i.e.' means 'that is' or 'in other words'.

were given a transparently bad reason (“Because I need to make some copies”). We often prefer bad or fabricated reasons to no reasons at all.

Outline of key points:

1. Unaware of relevant stimuli (causes)
  - (a) Ex: subliminal (implicit) perception
2. Unaware of our action (response)
3. Unaware of influence of stimulus on response
  - Ex: in many insufficient justification dissonance experiments and in attribution experiments, a large discrepancy between verbal reports of change and behavioral change
  - When asked, subjects were often very sure that the experimental manipulation had no effect on them (when it clearly did).

### Examples

.....

## 20.4 The Everyday Person as Intuitive Scientist

All of us, in our daily lives, have to solve many of the same sorts of problems that scientists do. Of course we don't go about it in such a careful way, but like scientists on the job, we have to form hypotheses, test them, give explanation, track down causes. Here are some of the things that both everyday people and scientists have to do.

1. Get data (sampling)
2. Draw Inferences from sample to population
3. Assess covariation (correlation)
4. Formulate Theories or Hypotheses
5. Test those Theories (theory maintenance and change)
6. Predict
7. Explain

These are similar to many tasks of actual scientists. How good are we at these things? The answers range from “Not very” to “O.k., but we could do a lot better.” We have already explored the first three items on this list, so we only need to recall them briefly here.

### 20.4.1 Gathering Data

Outline: There are several dangers at the very beginning of the process.

1. Selective Attention
2. Biases in perception (these affect classification of data)
  - Garbage in, garbage out

So data can be spotty and contaminated from outset.

1. Ex: stereotypes
2. Seeing what we expect to see
3. Not seeing what we don't expect to see

The key role of data to help us draw generalizations

#### Everyday Sampling

##### Inferences from Sample to Population

Good samples are not too small and they are not biased (good samples are representative) Bad inferences are often based on bad samples. This occurs, because we are often insensitive to the need for a large enough and unbiased sample.

.....

##### Insensitivity to Sample Size

Recall the two hospitals in Smudsville.

.....

##### Insensitivity to Sample Bias

Outline: Many causes of this. Here are a few examples.

1. Falling cats
2. Primacy and recency effects
3. Media exposure
  - Ex: causes of death – homicide vs. stomach cancer; pigs vs. sharks
4. Availability and salience (we often conduct sampling in our heads)
  - What is readily available in memory, imagination
  - Ex: More 6 letter words ending in 'ing' or with 'n' in fifth position?
  - Ex: (a) students read course evaluation from many students and (b) hear a few students make comments

- (b) has much more effect than (a)

.....

### Detecting Correlation (Covariation)

.....

#### 20.4.2 Testing and Predicting

Outline:

1. Black-box prediction (more of the same)
  - Trend extrapolation
  - Simple inductive inference: Kepler's laws and a new planet
  - Much forecasting
  - Correlations
2. Theory-driven prediction (something different)
  - Typically involves underlying causal mechanisms
  - Ex: Predictions about planets orbits from Newton's laws
3. Tradeoffs?

### Predictive Risks

Outline:

1. Tradeoff between detail and security
2. Problems if based on illusory correlation
3. Small or biased sample
4. Regression to the mean (we tend to overlook fact that high or low scores tend to be followed by more average ones).
  - (a) In test-retest situations bottom group on average improves and top group on average drops off
  - (b) Policies designed to counteract sudden rise in crime, disease,
  - (c) *Sports Illustrated Jinx*
5. Everyday (intuitive) vs. actuarial (algorithmic) predictions
6. Anchoring and Adjustment
  - Many of these things also problems at other spots

**Illusory Correlation Revisited** Outline: begin to enumerate **Illusory correlation**: thinking we perceive a correlation where one doesn't really exist

- Ex: Chapman gave subjects information allegedly about a group of mental patients (though in fact it was simply made up).
  - Everything fictitious and no actual correlations
  - Subjects “found” correlations between various diagnoses (e.g., paranoia) and features of drawings (e.g., weird eyes).
- Ex: Halo effect

Role of theory about the phenomenon

1. No theory: underestimate covariation
2. Prior theory: greatly overestimate covariation
3. Problematic: because when we “find” a correlation we often use it to make predictions and seek explanation for it.

Prediction and illusory correlation

.....

### Confirmation Bias

**Confirmation bias:** tendency to look for positive, but not for negative, evidence

Many studies (as well as a bit of careful observation) document our tendency to look for, and remember, and acknowledge the value of positive evidence that support our beliefs, while overlooking or undervaluing negative evidence that tells against them. This is called *confirmation bias* (Ch. 18).

Inconsistent pieces of evidence—inconsistent data—are often processed in a way that leads us to think that other people's behavior is more consistent than it actually is. If we formed a first impression that Wilbur is helpful and friendly, we are more likely to interpret his later angry outburst as justified by the circumstances or as an unusual lapse. There are limits to this, of course, and if Wilbur's later behavior is full of vicious and angry remarks, we will revise our earlier evaluation. But this will only occur after repeated pieces of data that disconfirm our original impression of Wilbur.

### Belief Perseveration

Once we form a belief or endorse a hypothesis, it can be difficult to dislodge, even when we get good evidence that the belief or hypothesis is false. In the chapters on memory we saw that this phenomenon is called **belief perseveration**. Once we arrive at a belief or hypothesis, whether it is accurate or not, it can be difficult to get rid of it.

### Stereotypes and Subtyping

Many of our stereotypes are also resistant to change, even after we meet numerous people in a group and discover that most of them don't fit our stereotype.

.....

### 20.4.3 Tracking Down Causes

.....

### 20.4.4 Giving Explanations

Attribution theory, and the fundamental attribution error.

**The Explanation Reflex.** We constantly seek to explain world around us; we are always on the look out for order, patterns, causes. The need to find explanations is usually a very good thing. It puts us in a much better position to predict what will happen under various conditions, which in turn helps us control our environment. But this tendency becomes problematic when we do not have enough information, or haven't thought about the issue carefully enough, to formulate a good explanation. Indeed, we like explanations so much that we often seem to feel that a bad explanation is better than no explanation at all.

The *Hawthorne effect*; the way people act changes when they know they are being studied. Perhaps this isn't terribly surprising, but it has important implications for studying human behavior in the real world. The effect gets its name because it emerged as an important problem in an early research project (beginning in the late 1920s) of Hawthorne Plant of the Western Electric Company in Cicero, Illinois. Researchers found that . . .

.....

## 20.5 Intuitive vs. Statistical Prediction

You are trying to predict who will win the various Big 10 basketball games this coming February. Here are three methods you might employ. This section is complete

1. Go with your own best educated guesses about who will win.
2. Go with the predications of experts (you can choose any experts you like, perhaps various coaches and sports writers).
3. Use a formula that has been carefully developed for making such predictions.

Most of us would expect (correctly) that the second method would work better (over the long run) than the first. What is surprising is that in many areas the third method works better (over the long run) than either the first *or* the second.

In fact “formulas” or computer programs have been found to make more accurate predictions than experts in many areas, to be just as good as experts in many others, and to be less accurate than experts in only a few. Among other things formulas are good at making predictions about medical or psychological health, about patients’ response to particular drugs or treatments, about success in school or on the job, about which prisoners are likely to violate parole, about which businesses are likely to go bankrupt, and so on.

**Where do the Formulas Come From?** Where do the formulas for such predictions come from? We could begin by thinking about the various sorts of evidence, the relevant factors or variables, that the experts think are good predictors for how teams will perform. These are called *predictor variables*. In our basketball example our initial list of predictor variables would probably include things like a team’s record, the record of their opponent, their margins of victory, how many starters are out with injuries, and so on. We would then write an equation (technically a “multiple linear regression equation,” but you don’t need to worry about that) or a computer program that takes these predictor variables as its input and returns predictions as an output.

At the beginning our formula and variables may not work all that well. The key is that we can adjust them in the light of experience, changing them as the season progresses. As we work on the program we might find that some variables we originally thought would tell us a lot (perhaps margins of victory) really don’t, while others we thought unimportant (perhaps records of opponents) turn out to be relevant. If we persevere, there is a good chance that before too long we will end up with a program that will do at least as well at predicting outcomes of games as any set of experts you choose.

Note that predictor variables needn’t disregard the human element. They may well include “subjective” things like the predictions of coaches; it is a completely empirical question which variables will lead to accurate predictions. The key point is that the program (which might or might not contain such “subjective variables”) is likely to do better than the experts at forecasting the outcomes of games.

**How do we Know the Formulas Work?** The early work on this topic was done by a clinical psychologist, Paul Meehl, in the mid 1950s. He marshaled a good deal of evidence to show that clinical psychologists and psychiatrists often did better to rely on the results of battery of psychological tests and other types of data than to rely solely on their own diagnoses (hence the issue is sometimes called “clinical vs. statistical prediction”).

Studies soon showed, however, that the same phenomenon occurred in many



other settings (e.g., decisions about parole, or about admission to graduate school). In all cases we see what predictions experts make, see what predictions the formulas make, then wait and *check to see* which predictions came true the most often. In study after study the formulas do at least as well, and often they do better.

Statistical predictions are based on correlations between the features of people or situations, on the one hand, and features of the phenomenon to be predicted, on the other. For example, computers using formulas that take account of college grades, scores on the LSAT, and a few other variables might be used to predict success in Law School. Perhaps the most surprising thing is that in many cases a rather simple equation with just a handful of predictor variables outperforms the experts. For example, four or five predictor variables might be enough to predict success in Law School at least as accurately as the judgment of law professors. This is often true even when the experts have access to all of the information in the predictor variables—and even when they know what the program or formula predicts.

**Why does Statistical Prediction Work?** When we stop to think about prediction, it shouldn't be too surprising that formulas often outperform humans. When we predict which team will win a baseball game, we need to integrate the information from a number of variables (records thus far, records of the starting pitchers, records of opponents, players out with injuries, and doubtless a few other things). The problem lies in combining these diverse pieces of information in a way that will lead to accurate predictions.

We do not, in general, think that we are particularly skilled at combining diverse pieces of information. As Meehl notes, if you have just put a number of things in your grocery cart without keeping track of their prices, you don't think that you can simply glance at the cart and come up with a more accurate estimate of the total price than the one the cash register rings up.

There are many reasons why cash register addition is more accurate than intuitive estimation. They include our selective attention, limited working memory and computational capacity, and susceptibility to various biases (e.g., primacy and recency effects). And the situation is even more complicated in predicting which team will win a game or which people will flourish in a given job. Here all of the limitations just noted are present, and more easily come into play, including confirmation bias, susceptibility to illusory correlations, underestimation of regression effects, failure to adjust for extreme anchors, overreliance on the availability heuristic, susceptibility to the dilution effect, not to mention wishful thinking.

**Real-Life Implications** Even if a formula would be better, where would we find one? Few of us would hole up over Spring Break to write a program to make predictions, and life is short enough that we shouldn't, even if we knew how. But issues about statistical vs. intuitive prediction are important in trying to decide whether various policy procedures (e.g., for admitting people to a college or determining which businesses are good bets for loans) are good ones, and in many cases decision makers have access to a great deal of data and a good formula that would enable them to make better predictions than many of them actually make.

Furthermore, even where a formula hasn't been devised, such agencies have large budgets, and it would often be cheaper in the long run to develop an accurate formula than to continue to make less than optimal predictions. Closer to home, you might someday serve on a hospital board, a committee to make recommendations about admissions to schools or on a parole board, and it is worth remembering that your gut reactions are not likely to be all that accurate.

Even closer to home, tables and programs have been developed for making predictions about various medical treatments for illnesses that you (or a loved-one) may eventually have. At the very least, the points noted in this section are a good reminder that we shouldn't simply accept an expert's predictions, or their claim that their method works better than more mechanical alternatives. It is an empirical question whether this is true or not, and in many cases the track records of experts are not as good as they suppose.

It is worth repeating, since the point is sometimes missed, that nothing we have said rules out including subjective judgments among the predictor variables. A clinical psychologist's interview with a patient or a personnel officer's interview with a job candidate may supply highly relevant information. Whether that is so is an empirical matter, and if it does, it should be included among the predictor variables.

The much greater danger, though, is that such subjective judgments often lead us to completely ignore other information that is at least as, or even more, relevant. We should note one danger that statistical prediction share with many other quasi-formal procedures; it can (though it need not) encourage us to focus on things that are easily measured or quantified while ignoring other, perhaps more important, factors.

**Not Every Case is a Broken-Leg Case** Wilbur believes that statistical prediction is more accurate than intuitive prediction when it comes to basketball games, and so he places his bets on the teams a formula favors. The starting center for the Tigers, the team he plans to bet on, has pretty much been carrying them all season, and just as Wilbur is about to phone his bookie, he learns that she broke her leg and

won't play.

What should Wilbur do? The answer is obvious; he should use commonsense and not bet on the Tigers. *Sometimes* there will be *exceptions*, cases where the formula very obviously gets things wrong. When this occurs, we should ignore the formula. But the vital point here is this: *most cases are not exceptions*. Neither the formula nor the experts are infallible, and an expert will (over the long run) get some cases right that the formula gets wrong. But in areas where a formula does better, overall, than the experts, *there will be even more cases that the formula gets right and that the expert gets wrong*. This is just simple arithmetic.

It is natural for experts to focus on those cases where they were right and the formula was wrong. This may result from various self-serving biases, but it need not. It could, for example, result in large part from confirmation bias (page 359), from the common tendency to notice and remember confirming or positive evidence, while ignoring disconfirming or negative evidence. But if we want an accurate picture of track records, we must also consider the negative cases, those where the formula was right and the expert was wrong.

This is just one of many sorts of situations where it can be tempting to treat a large number of cases as exceptions (“I know that as a general rule we should expect such and such, but this case is special . . .”). But by definition it is impossible for most cases to be exceptions.

It is impossible for most cases to be exceptions

It is true that many of us trust our intuitive judgments, our gut reactions, more than statistical or actuarial methods. After all, how can a computer or a formula consider the human factor; how can it use commonsense when it's needed? Computers aren't good at these things, and this does sometimes lead them to make poor predictions. But unlike people, they don't have a tendency to treat all, or even most, cases as exceptions, as broken-leg cases. Moreover, they are immune to the various limitations and biases and fallacies that mar human reasoning. And in cases where the statistical methods are better, our confidence in intuition or commonsense is often misplaced.

### Exercises

1. Explain how some of the biases and fallacies we have studied could lead people to place more weight on a personal interview (e.g., with a job applicant or a prospective patient) than they should.
2. You have just been hired to oversee admissions to one of the smaller colleges at your school. You have a small group of people to help you, but the task isn't easy since you can only admit about one in six applicants. What sorts of things would you use in deciding who gets admitted and what would you look for? How would

the final decision be made?

3. What are some of the main variables that you think would be good at predicting how likely someone would be to violate parole? Explain in detail how you would determine whether these variables really did lead to accurate predictions.

4. It is sometimes objected that using statistical prediction dehumanizes things. One student who read an earlier version of this section made the point well: “I know statistical prediction is supposed to be more accurate at predicting who will do well at a given job, but you are hiring a person, not a number.” It is natural to feel some sympathy with this view, but if statistical prediction *is* more accurate in a given domain, is it better, or more humane, to go with human intuitive prediction instead? Give an example and present your answer in terms of it.

### We Aren’t Merely Intuitive Scientists

Of course people have many goals besides consistency and accuracy. Human beings are not merely intuitive or everyday scientists or intuitive statisticians: we are also intuitive lobbyists and defense attorneys, intuitive conflict negotiators and counselors, intuitive gossip columnists and romance novelists, intuitive televangelists and inquisitors. The spirit of this criticism is right; it often makes sense to violate the prescriptions of normative models when we have additional goals like drawing quick conclusions, or avoiding decisions that could lead to regret, disappointment, guilt or embarrassment, or making choices that will minimize conflict or allow deniability. Nevertheless, . . .

## 20.6 Pseudoscience

According to a Gallup (Gallup News Service) poll result reported on June 8, 2001, belief in various paranormal phenomena increased during the 1990s. Over half of those polled believe in extrasensory perception (ESP), and at least one in three believe in ghosts, haunted houses, telepathy, clairvoyance, devil possession, or that extraterrestrial beings have visited the earth. In this section we will examine some of these beliefs and their connections to the things we have studied thus far.

Things like the astrology column in your local newspaper are typically harmless fun. But other pseudosciences, e.g., fraudulent medical or psychological treatments and worthless medical treatments (this includes a fair bit of holistic medicine) cost people vast sums of money and untold misery. We will use our study of genuine science and of various causes of defective thinking as a background for an examination of pseudoscience.

This section under construction

There are several features that are pretty good symptoms of a pseudoscience: *untestability, use of vague predictions, use of multiple predictions, saliency of successful predictions, fraud, self-fulfilling prophecy, wishful thinking, the P. T. Barnum effect, dissonance reduction, and true believers*. We have met some of these notions before and some we will study in more detail in later chapter.

One way to get a feel for the difference between science and pseudoscience is to look at some contrasts: astronomy vs. astrology; parapsychology vs. modern experimental psychology, . . .

1. We saw above that a theory or hypothesis is **testable** just in case some sort of objective, empirical test could give us some reason to think that it is true or else some reason to think that it is false. So a theory is *untestable* if this cannot be done. Some pseudosciences are not testable, and so aren't even in the running as scientific theories. However, many pseudosciences *can* be tested. It's just that they fail the tests, often rather badly, or that they are never subjected to *severe* tests because their practitioners make extremely vague predictions or predictions that we all knew could come true quite independently of the pseudoscience.
2. When practitioners of a pseudoscience bother to make predictions at all, they usually make very **vague predictions**. Their predictions are so unspecific that they are likely to be true of almost anybody (e.g., "something important will happen in your life in the next year," "you like to help others"). It is then easy for a person who is sympathetic to the pseudoscience to think that the prediction applies to them.

When a prediction is sufficiently vague, it provides a very weak test of a theory. No matter how things turn out, it will be difficult to show that the prediction didn't come true (since the predictor can always "reinterpret" it after the fact in such a way as to make it seem correct).

For example, the "Astro-graph" in the *Norman Transcript* recently said of people with my astrological sign: "You will do well to cooperate with others today. You are as much a giver as you are a taker. Your attitude inspires cooperation." This is vague enough to apply to a great many people (in general, most people would do well to cooperate with others). Because of their vagueness, most predictions of pseudosciences do not provide good tests of their hypotheses.

The characterization from the Astro-graph also contains another symptom of pseudoscience, the use of a flattering characterization. Human nature being

what it is, most people are primed to believe flattering descriptions (how many people really think that they aren't capable of working with others?).

3. Many practitioners of pseudosciences make **multiple predictions**. For example, an astrologer will often make four or five predictions. This increases the likelihood that one or two of the predictions will come true, and because of the saliency of successful predictions, these are likely to be remembered, while the (more numerous) unsuccessful predictions are forgotten.
4. The **saliency of successful predictions** means that predictions that come true tend to be remembered, whereas predictions that don't are not. If enough predictions are made, chances are good that a few will come true, and since these are the ones that will tend to be remembered (especially since these are the ones the predictor will remind us about), even hopeless hypotheses may seem to be confirmed.
5. Although much pseudoscience is practiced by people with good intentions, some is the result of deliberate **fraud**. Various pseudosciences, e.g., fraudulent medical treatments, can be highly profitable, since there is always a quick buck to be made from some gullible person. The fact that many pseudosciences are a good business provides a strong motive for con artists masquerading as scientists to push their claims.
6. A **self-fulfilling prophecy** (Module 9) is the tendency for an individual's expectations about the future to influence that future. Pseudosciences sometimes involve self-fulfilling prophecies; for example, if an astrologer tells you that you will do a certain thing today (e.g., make a special effort to cooperate with others), it may make you more likely to do just that.
7. The **P. T. Barnum effect** is named after the founder of the world-famous circus, who said "there's a sucker born every minute." There is a little gullibility in the best of us. Some part of us is captivated by the mysterious and charmed by novel and unexpected claims, and ready to accept them with little evidence. After all, pseudosciences and superstitions are more fun than the same dull old facts. And in those cases where a pseudoscience promises us something that real science can't deliver (a prediction about our love life, a cure for a deadly disease), it is not surprising that many people *want* the claims to be true (which in turn makes them more disposed to believe the claims).
8. **Wishful thinking** occurs when the desire that some claim be true leads us to believe that it *is* true. But we are all susceptible, and in its more minor forms

it is quite common. Our tendency to wishful thinking is one reason why claims by pseudoscientists, advertisers, and many others are taken to be true even though there is little evidence in their favor. For example we would like a quick fix for many of our problems, and although fast and easy solutions are often too good to be true, we would like to think that they would work.

9. A **true believer** in a theory or a movement is a person so deeply committed to it that he interprets everything in terms of its concepts and slogans. True believers in a cause are especially prone to wishful thinking. Indeed, they will do almost anything to avoid acknowledging the shortcomings of their pet theories. Some of the maneuvers used for this involve intentional deception (and so involve some of the concepts studied in module 8 on advertising and PR). But sometimes the true believer deceives himself or relies on a defense mechanism (we'll learn more about this in module 10, when we discuss the effect of emotions and wishful thinking on reasoning).

### **Pseudoscience and Coincidence**

This is yet another situation in which confirmation bias (page 359), the common tendency to notice and remember confirming or positive evidence, while ignoring disconfirming or negative evidence, can lead to a distorted picture of things. We may remember cases where an astrologer or someone down at the psychic hotline was right about something; we may forget all the cases where they were wrong.

### **Quacks, Scams, Snake Oil, and Fraud**

.....

## **20.7 Chapter Exercises**

1. Find an example of a pseudoscientific claim or prediction in something you read outside class, and bring a copy and a one paragraph analysis of it to class. Explain in your own words why it is pseudo, rather than a genuine, scientific claim.
2. Make three predictions (in about a sentence each) that are so vague that it would be very difficult to show them false. Then in a second sentence say why the prediction would be difficult to disprove.
3. Make three predictions on the same topic that differ in the severity of test they would probably provide for a theory that made them. Explain which are most severe, and why.



4. Bring a horoscope to class that covers the week just before this assignment is to be done (it may be from a newspaper, a magazine, or, if you know an astrologer, have one be cast for someone who was born on noon, April 7, 1975 in Norman). We will then consider a few signs to see how accurate the horoscopes were.
5. Give an example of a self-fulfilling prophecy that is based on some superstition or pseudoscientific claim.<sup>2</sup>

---

<sup>2</sup>Carl Hempel has a very clear discussion of Semmelweis's discovery of the cause of childbed fever in his *Philosophy of Natural Science*, Prentice-Hall, 1966, Ch. 1. Although it's not easy reading one of the best, and certainly the most entertaining discussion of significance testing is Paul E. Meehl, "Theoretical Risks and Tabular Asterisks: Sir Karl, Sir Ronald, and the Slow Progress of Soft Psychology," in his *Selected Philosophical and Methodological Papers*, U. of Minnesota Press, 1991. Lee Ross, "The Intuitive Psychologist and his Shortcomings: Distortions in the Attribution Process," in Leonard Berkowitz, ed., *Advances in Experimental Social Psychology*, Volume 10 (New York: Academic Press, 1977). R. E. Nisbett & T. D. Wilson "Telling More than We Can Know: Verbal Reports on Mental Processes," *Psychological Review*, 84, 1974, 231-259. The phrase "voodoo science" is Robert L. Park's, and he discusses a number of recent examples in *Voodoo Science: The Road from Foolishness to Fraud*, Oxford University Press, 2001. A good, up-to-date discussion of the relative accuracy of informal (intuitive) vs. formal (statistical) methods of prediction may be found in William Grove and Paul Meehl, "Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical Statistical Controversy," *Psychology, Public Policy, and Law*, 2, 1996; 293-323. A thorough overview may be found in Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (forthcoming). *Clinical vs. Mechanical Prediction: A Meta-analysis*. *Psychological Assessment*. Further references here will be supplied when chapter is finished.



## 20.8 Appendix: Scientific Notation and Exponential Growth

The mathematics of uncontrolled growth are frightening. A single cell of the bacterium *E. coli* would, under ideal circumstances, divide every twenty minutes. That is not particularly disturbing until you think about it, but the fact is that bacteria multiply geometrically: one becomes two, two become four, four become eight, and so on. In this way it can be shown that in a single day, one cell of *E. coli* could produce a super-colony equal in size and weight to the entire planet Earth.

— Michael Crichton, *The Andromeda Strain*, 1969, N.Y.: Dell, p. 247.

A billion here, a billion there, and pretty soon you're talking real money"  
— often attributed (probably inaccurately) to Senator Everett Dirksen  
during Congressional Budget Debate

### 20.8.1 Scientific Notation

*Scientific notation* affords a compact way of representing numbers, especially those that are very large or very small. When you encounter numbers like 123, 839, 203 or 0.0007073 it is difficult to form any clear idea what they mean. Scientific notation is easier to read, allows us to tell at a glance what the order of magnitude is (which, unless we are engineers or scientists, is all we typically need to worry about), and makes it easier to use the number in calculations. The basic idea behind scientific notation is not difficult, and you don't need to know much about science to understand it.

This appendix is complete

#### How it Works

Here are two examples of scientific notation:  $5.33 \times 10^2$  and  $9.31 \times 10^{-3}$ . Like all numbers in scientific notation, they consist of two parts: first a number at least as big as 1 but less than 10 and, second, a power of ten. Lets consider these two features in turn.

The first part of the representation is a number greater than 1 but less than 10. What this means is that in scientific notation, the decimal point always occurs immediately to the right of the first non-zero digit. So 1, 3.33, 7.02 and 9.99 are allowed. The following, however, are not allowed: .5 (it's less than 1), 32 (it's not less than 10), or 10 (it's also not less than 10).

The second part of a number in scientific notation is a power of 10, i.e., 10 with some exponent (which is a superscript). This means that it is 10 raised to some power. By convention  $10^0$  (as with every number raised to the zero-th power) is 1 (a 1 with no 0s). Working our way up,  $10^1 = 10$  (a 1 followed by one 0),  $10^2 = 100$  (a 1 followed by two 0s),  $10^3 = 1000$ , and so on up. We can also work our way down:  $10^{-1} = 0.1$ ,  $10^{-2} = 0.01$ ,  $10^{-3} = 0.001$ , and so on.




Figure 20.3: Powers of Ten

**Numbers Greater than 1** We begin by seeing how to translate standard numerical notation for numbers greater than 1 into scientific notation. The number 300 is  $3 \times 100$ , so we can write it as  $3 \times 10^2$ . And 5000 is  $5 \times 1000$ , so we can write it as  $5 \times 10^3$ . What about 320? This is  $(3 \times 100) + (2 \times 10)$ , so we can write it as  $3.2 \times 10^2$  (since  $.2 \times 10^2 = 2 \times 10$ ). The first number (in this case 3.2) is called the *base* or *mantissa*, and the second number, the power of ten, is called the *exponent*. On some hand calculators and a few computer programs,  $2.54 \times 10^{11}$  might be expressed as 2.54E + 11 (E for exponent).

There is a very simple rule that applies to all numbers greater than 0. We raise the mantissa to the number of places in the original number after the first digit. If the number has  $n$  digits, then we will use 10 to the  $n - 1$ . Examples should make this clear.

- 500,000 has five digits (in this case all 0s) after the first digit (here a 5). So we represent it as  $5 \times 10^5$ .
- 27,300 has four digits after the first digit (here a 2). So we represent it as  $2.73 \times 10^4$
- $429,3120 = 4.2931 \times 10^6$
- $84,579,800,000 = 8.45798 \times 10^{10}$

We pronounce  $3.26 \times 10^3$  as “three point two six times ten to the third.”

**Numbers Less than 1** Scientific notation is also useful for numbers between 1 and 0, especially those that are very small. Remember that 0.1 is  $\frac{1}{10}$ , 0.01 is  $\frac{1}{100}$ , 0.001 is  $\frac{1}{1000}$ , and so on. We can connect this to exponents this way:  $10^{-2} = \frac{1}{10^2}$  or  $\frac{1}{100}$ , which is one one-hundredth, or 0.001. Similarly,  $10^{-3} = \frac{1}{10^3} = \frac{1}{1000}$ , which is

0.001. When the number is less than one, we use a power that is a negative number. The pattern here is

- $10^{-1} = 0.1$
- $10^{-2} = 0.01$
- $10^{-3} = 0.001$
- $10^{-4} = 0.0001$

The power is a negative number equal to the number of zeroes *after* the decimal point and *before* the first non-zero digit, plus one. Equivalently, it is the number of zeroes *after* the decimal point up to and including the first non-zero digit. This sounds much more complicated than it actually is, once we look at a few examples.

- $0.124 = 1.24 \times 10^{-1}$
- $0.00054 = 5.4 \times 10^{-4}$  (it has 4 – 1 zeros after the decimal point, so the exponent is 4)
- $0.000000375 = 3.75 \times 10^{-7}$

There are several standard prefixes for referring to common powers of ten.

Number	Prefix	1 / Number	Prefix
One thousand	kilo-	one one-thousandth	milli-
One million	mega-	one one-millionth	micro-
One billion	giga-	one one-billionth	nano-
One trillion	tera-	one one-trillionth	pico-

For example, the size of hard drives on new computers is typically measured in gigabytes, or billions of bytes.

It is also easier to perform operations on numbers when they are expressed in scientific notations. For example, to calculate  $(3 \times 10^2) \times (2 \times 10^4)$  we multiply the two  $(3 \times 2)$  and *add* the exponents. Unless we are scientists or engineers, we will often only be interested in ball parks figures, so we can forget about the mantissas and simply add the exponents to get an order of magnitude. Adding exponents in this case gives us  $10^6$  on a million. And to divide two numbers expressed in scientific notation we subtract exponents:  $3 \times 10^7 / 2 \times 10^3$  leaves us with an order of magnitude of  $10^4$ , or ten thousand.

### Exercises

1. Express the following numbers in scientific notation:

1. 47,392

2. 123,000,000,000
3. 0.00743
4.  $\frac{3}{10,000}$

Translate the following numbers from scientific notation to their more conventional representation.

1.  $5.34 \times 10^4$
2.  $5.34 \times 10^7$
3.  $5.34 \times 10^{-7}$
4.  $3.22 \times 10^{-2}$

### 20.8.2 Exponential Growth

We keep adding 1 to the exponent of 10 in the series:  $10^2, 10^3, 10^4, \dots$ . Such sequences are said to grow *exponentially* or *geometrically*. The rate of growth in each succeeding step is proportional to the size of the number in the preceding step. It often comes as a surprise how huge such sequences become after even a modest number of steps. People tend to underestimate how rapidly such sequences grow, but in fact the growth is very explosive, and such progressions are sometimes said to “blow up.”

One way to get a feel for just how overwhelming exponential growth can be it to plot a few numbers on a calculator (you will be asked to do that in the exercises). There are also various analogies that help us get some grip on the idea. One thousand seconds lasts about 17 minutes. Now lets compare this to the length of time various other numbers of seconds would take to elapse (we’ll round off, since we are only interested in the rough size of the differences).

Name	Numeral	Power of 10	Length of that many seconds
One thousand	1000	$10^3$	17 minutes
One million	1,000,000	$10^6$	11.5 days
One billion	1,000,000,000	$10^9$	32 years
One trillion	1,000,000,000,000	$10^{12}$	32,000 years

Here is another standard example that also highlights the explosive nature of exponential growth. A chess board or check board is an eight by eight square with 64 positions. Half are light and half are dark, but we’ll leave them all the same color in Figure 20.4 on the next page for legibility.

We begin by putting two cents on the upper left square, twice that amount in the square to its right, and so on across the board, doubling the amount as we go. By the last square the amount of money would be about  $\$1.85 \times 10^{17}$ . The growth

2¢	4¢	8¢	16¢	32¢	64¢	\$1.28	\$2.56
\$5.12	\$10.24	\$20.48	\$40.96	\$81.92	\$163.84	\$327.68	\$655.36
\$1310.72	\$2621.44	...					

$2^{64} \text{¢} = 18446744073709551616 \text{¢}$   
 $= \$184467440737095516.16$   
 $\approx \$1.84467 \times 10^{17}$   
 —which is way over one trillion ( $10^{12}$ ) dollars

Figure 20.4: Doubling the Money on Squares

$$\begin{aligned}
 2^2 &= 4 \\
 2^3 &= 8 \\
 2^{10} &= 1024 \\
 2^{20} &= 1048576 \\
 2^{30} &= 1,073,741,824 \\
 2^{50} &= 1,125,899,906,842,624 \\
 2^{100} &= 1,267,650,600,228,229,401,496,703,205,376 \\
 2^{150} &= 1,427,247,692,705,959,881,058,285,969,449,495,136,382,746,624
 \end{aligned}$$

Table 20.1: Exponential Growth – Powers of Two

begins relatively slowly, but before long it becomes overwhelming. Table 20.1 gives another representation of how dramatic exponential growth can be. When you deposit money in a savings account that pays compound interest, your money grows exponentially. If the interest rate were 7% a year, the money would double in ten years. Here ten years is said to be the *doubling time*. This is the number of years (or some other unit of time) it takes some quantity like money or population to double.

**Doubling time:** the number of years it takes for a quantity growing exponentially to double in size

Exponential growth also applies to populations, and here the results are often less happy. Just as with money in the bank, if a population grows at 7% a year it will double in ten years. Suppose that Belleville grew at this rate for the last several decades. If it had a population of 10,000 in 1972, that would have doubled to 20,000 in 1982, doubled again to 40,000 in 1992, and would be 80,000 now. With exponential growth increases often seem slight at the beginning, but when they begin to get large, they get very large very quickly.

It is worth learning to calculate things like doubling time exactly, but for many purposes approximate figures are good enough. One useful rule technique is the so-called *rule of 70*.

**Rule of 70:** to find the approximate doubling time of a variable that grows at a given percentage a year divide the percentage into 70.

To make our examples easy to work with we assumed a growth rate of 7% a year. When we divide 7 into 70 we get 10. If the growth rate were 10% then (since  $70/10 = 7$ ), the doubling time would be about 7 years. As a final illustration, a growth rate of 2% would yield a doubling time of about 35 years.

**Exercises**

1. Fill out the rest of the third row of the chess board in Figure 20.4 on page 423. How far can you go calculating the numbers for the squares on your calculator?
2. Smudsville has had a constant growth rate of 7% for the last five decades. It's population in 1952 was 30,000. What, approximately, is its population now? Explain how you arrived at your answer.
3. Wilburtown has had a constant growth rate of 5% for the last five decades. It's population in 1952 was 30,000. What, approximately, is its population now? Explain how you arrived at your answer.





# Chapter 21

## Risk

**Overview:** In this chapter we study risks, the misperception of risks, and ways to more accurately assess the riskiness of various actions and projects.

### Contents

---

<b>21.1 Life is Full of Risks</b> . . . . .	<b>428</b>
<b>21.2 Describing Risks</b> . . . . .	<b>429</b>
21.2.1 Risk Ratios . . . . .	429
21.2.2 Exercises . . . . .	431
21.2.3 Finding Information about Risks . . . . .	433
<b>21.3 Health Risks</b> . . . . .	<b>433</b>
21.3.1 The Big Three . . . . .	434
<b>21.4 Crime Risks</b> . . . . .	<b>436</b>
21.4.1 Exercises . . . . .	437
<b>21.5 Other Risks</b> . . . . .	<b>437</b>
21.5.1 Sex . . . . .	437
21.5.2 Love and Marriage . . . . .	437
21.5.3 Jobs and Businesses . . . . .	437
<b>21.6 Cognitive Biases and the Misperception of Risk</b> . . . . .	<b>437</b>
21.6.1 Tradeoffs . . . . .	441
<b>21.7 Psychological Influences on Risk Assessment</b> . . . . .	<b>441</b>
21.7.1 Individuals Differences . . . . .	441
21.7.2 Groups . . . . .	442

---

**21.8 Chapter Exercises** . . . . . 442

---

## 21.1 Life is Full of Risks

There is no way to avoid all risks

Nothing will be quite the same after September 11, and that certainly includes how we think about risk. One moment things are fine, the next moment disaster strikes. Indeed, I was beginning the first draft of this chapter on May 3, 1999, when a devastating tornado struck Oklahoma City. It killed over forty people, caused hundreds of injuries, and did millions of dollars worth of damage to property.

It was a sobering reminder of how risky life can be and of how little, sometimes, we can do to avoid those risks. But these tragic consequences also emphasize the importance of doing what we can to avoid life's dangers. In this chapter we will learn how to distinguish the big risks from the small ones, and we will develop some tools for thinking about risks.

*Everything* involves some risk. Even if you stay home in bed with the covers pulled up tight accidents can still befall you. Every year thousands of people are taken to emergency rooms after falling out of bed. Some risks are serious and we should take precautions to avoid them. Other risks are overblown, and we just make ourselves miserable if we dwell on them. The trick is to learn how to tell the difference.

Before proceeding, take the following pretest. Indicate which of the two items on each row you think causes more deaths in America each year; you might also indicate how much more likely you think one is than another (answers are given at the end of the chapter).

### Pretest

- |                      |                   |
|----------------------|-------------------|
| 1. 1. Diabetes       | 2. Homicide       |
| 2. 1. Suicide        | 2. Homicide       |
| 3. 1. Asthma         | 2. Tornado        |
| 4. 1. Lightning      | 2. Flood          |
| 5. 1. Lung Cancer    | 2. Heart Disease  |
| 6. 1. Stomach Cancer | 2. Rabies         |
| 7. 1. Homicide       | 2. Stomach Cancer |
| 8. 1. Suicide        | 2. Syphilis       |

There are many sorts of risks: health risks (“Is rabies really a danger?”), physical risks (“Is skydiving more dangerous than rock climbing?”) job risks (“What

are the chances that a new restaurant will go under its first year?”), financial risks (“What if I buy stock in a company that goes bankrupt), crime risks (“What are the chances my car will be stolen?”), social risks (“Will I be a social outcast if I tell people what I really think about gun control?”), sexual risks (“How reliable are condoms?”), environmental risks (“How real *is* the threat of global warming?”). You name it, and there’s a risk involved.

We will touch on various sorts of risks, but to keep things manageable we will focus on causes of death and types of crime, with a bit on some other risks you encounter frequently. But the tools we develop for thinking about these types of risks apply equally to all of the other types.

## 21.2 Describing Risks

We understand risks better when we can describe them in rough numerical terms (e.g., about one person out of every forty two is killed in an automobile accident). Precise numbers don’t really matter; ball park figures are enough. It won’t matter to most people whether 37,000 people out of a hundred million or 43,000 people out of a hundred million die each year from lung cancer (it’s the former). But it does make a difference with it’s 37,000 people out of a hundred million or 370 out of a hundred million.

It will make things more realistic if we work with some actual numbers, so we will use Table 21.1 on the following page, which gives the leading causes of death among all Americans in 1997. These figures are based on a report released by the National Center for Health Statistics (the number after each cause of death is the actual number who died; data are based on a review of death certificates).

### 21.2.1 Risk Ratios

Risks are reported by fractions. They are numbers from 0 to 1, and they can be interpreted as *probabilities*. We will call these *risk ratios*. In the case of death rates, the risk ratio is given by a fraction:

$$\frac{\text{Number of Deaths}}{\text{Number in Target Population}}$$

In the case of deaths the numerator is clear cut; it is simply the number of people who died from a given cause. But the denominator is less clear cut, and in many cases there will be different ways to express it. For example, in assessing the risk of hang gliding we would want to express the number of deaths per (over) people who went hang gliding or the number of hours spent hang gliding. We will return to this

1. Heart disease: 725,790
2. Cancer: 537,390
3. Stroke: 159,877
4. Lung diseases: 110,637
5. Accidents: 92,191
6. Pneumonia and influenza: 88,383.
7. Diabetes: 62,332
8. Suicide: 29,725
9. Kidney disease: 25,570
10. Liver disease: 24,765
11. Blood poisoning: 22,604
12. Alzheimer's disease, 22,527
13. Homicide: 18,774
14. HIV and AIDS: 16,685
15. Hardening of the arteries: 15,884
16. All other causes: 361,635

Table 21.1: Causes of Death in America in 1997

important point below. But in the present case we are dealing with conditions that could strike almost anyone, so to keep things simple we will use the total number of Americans as our denominator.

In the 1990 census the number of Americans tallied just under 250 million. This figure is low, since several million people weren't counted and the population has risen since then. But 250 million is a nice round number, so we'll use it as our count for the total number of Americans. So we express the *death rate* for a given medical condition as

$$\frac{\text{Number of Deaths from Condition}}{\text{Total Number of Americans}}$$

which is

$$\frac{\text{Number of Deaths from Condition}}{250,000,000}$$

For example, in 1997 725,790 people died from heart disease, so the death rate from heart disease is:

$$\frac{725,790}{250,000,000}$$

Numbers with such large denominators are hard to comprehend. It is possible to round such fractions off and reduce them down to more meaningful numbers, but this can take time. We can also use a calculator to divide the numerator by the denominator. This gives a decimal value which is in fact the probability or frequency of a heart attack death. In 1997 this probability number is 0.0029. But this number is so small that it's hard to comprehend. We need a more user-friendly way to express these numbers.

### Deaths per Million

It is often clearest to express risk statistics in terms of number of deaths per million, per hundred thousand or the like. This also makes it easier to compare risks. The general formula for this is just:

$$\frac{\text{Number of Deaths}}{\text{Number in Target Population}} \times c$$

where  $c$  is the common denominator we use for all of the risks. For example, if we want to express the risk ratio as number of deaths per million, then  $c$  is one million (1,000,000). If we want to express them as number of deaths per hundred thousand, then  $c$  is one hundred thousand (100,000).

The choice of  $c$  is a matter of convenience: select a number that will make the resulting figures as easy to understand as you can. When we are talking about the entire United States, thinking in terms of deaths per one million or even per hundred million makes sense. But if we were thinking about death rates in Norman, a smaller number like deaths per one thousand would be easier to understand.

The number of heart attack deaths per million people is

$$\frac{725,790}{250,000,000} \times 1,000,000,$$

which is approximately 2904 per one million people. In 1997, 2904 people out of every one million had a fatal heart attack.

#### 21.2.2 Exercises

1. Express the death rates for cancer, stroke, diabetes, and suicide
  1. as a fraction
  2. as a probability
  3. in terms of number of deaths per million people
  4. in terms of number of deaths per hundred million people

### Finding a Useful Denominator

Since you probably won't be compiling risk tables anytime soon, you won't have to make decisions about the best denominator to use. But you do need to think about the issue, so that you can more easily interpret figures that you read.

If we are thinking about the relative risks of rock climbing and driving a car it won't be very useful to express them in terms of the number of deaths (or injuries) out of all 250 million people living in the United States. If we do this driving a car will look much more risky, since so many more people drive. Instead we want a denominator that reflects just the actual number of people involved in each activity. We could use the number of people who go rock climbing each year, so that our ratio is number of injuries, say, per number of rock climbers. We could also express things in terms of the number of hours spent rock climbing each year, so that the figure is number of injuries over the total number of hours people spent rock climbing in a given year.

For example, if 700,000 people went rock climbing last year and 150 of them were killed, the relevant ratio would be  $\frac{150}{700,000} = \frac{15}{70,000}$ . Alternatively we could use the number of total that people hours spent rock climbing in 1997. Suppose that it is 900,000. If we do this we would end up with a figure that told us about the number of deaths per 900,000 hours. We could also use a common denominator figure like we did above to the number of deaths per 100,000 hours (or 1000 hours, or whatever makes the most sense). It's useful to try to understand these general ideas, but you won't need to worry about the details of all of this.

When we think about the riskiness of an occupation (e.g., coal mining), we will probably begin with number of deaths per year over the number of people who were coal miners in that year. But these yearly risks are cumulative. So if we want to know the risk confronting someone who spends their entire working life as a coal miner, we need to multiply this figure by the number of years the average person works (about 40). This gives the risk that a coal miner will die in the mines at some point in her career.

Choosing an informative denominator is largely a matter of commonsense. For example, suppose we are planning a trip to Anchorage, Alaska and we are wondering about the relative risk of driving or taking an airline. We could look at the number of deaths per hour for driving and for flying. But since we are interested in the relative risk of driving or flying for the entire trip, it makes more sense to look at the numbers of deaths per mile. After all, we have to travel about the same number of miles with either mode of transportation.

### 21.2.3 Finding Information about Risks

The United States Census Bureau (<http://www.census.gov/>) has an extensive web site with a huge collection of data (so much that it can be difficult to find what you want). Another useful source is the National Safety Council's Environmental Health Center (<http://www.nsc.org/ehc.htm>). But for specific topics it is often easiest to log on to the net, fire up your favorite search engine, and search for the specific information you want. Some of it is reliable, some of it isn't; as always, the checklist (p. 120) for evaluating information on the web is relevant here. All of the figures cited in this chapter are either from the net or from the 1999 *New York Times Almanac*.

## 21.3 Health Risks

In general people tend to *underestimate high* probabilities and *overestimate low* probabilities, and this holds in the area of risk assessment. You hear about rabies shots and cholera shots. These are both serious diseases, and it makes sense to get a rabies shot for your dog or a cholera shot for yourself if you are traveling to countries where it is a danger. But in 1997 only two (human) Americans got rabies and only seven were stricken by cholera. So isn't a good use of your time thinking about other ways to avoid rabies or cholera (not to mention the Plague, which sounds terrible but killed just four Americans in 1997).

It is true that anyone can be stricken by a rare disease, and it won't be much comfort for that person to hear that their probability of getting it was low. But *everything* carries some risks, and we simply can't plan to deal with every risk life has to offer. The best approach involves two steps, both of which involve probabilities.

1. Identify the things that are large risks in your life (or your family's) life.
2. Adopt the measures that have the highest probability of helping you avoid those risks.

We have consider the things that bear on the second issue in earlier chapters, so here we will focus on the first. But one general point is worth emphasizing before turning to details. Whenever people are worried about a risk, someone will come along with a way for you to avoid or reduce it—for a price. If their remedy sounds too good to be true, it probably is. On the other hand, we do know a good deal about how to reduce many of life's most serious risks, and often the ways to do it don't require you to spend any money at all.

### 21.3.1 The Big Three

Table 21.1 on page 430 on death rates in 1997 gives a good indication of the causes of death in America (and in Oklahoma) over the recent past (although deaths from AIDS rose rapidly in recent years, it is now declining sharply and is no longer one of the top ten causes of death). As the table shows, there are three big killers: heart disease, cancer, and stroke. Your chances of being killed by one of these is *much, much* greater than your chances of dying in any sort of accident, or from anything else. Fortunately, there are also reasonably simple steps you can take to greatly lower your risk from the Big Three.

#### Heart Attacks

A heart attack occurs when the blood supply to the heart is blocked. About 1.5 million Americans suffer heart attacks each year. In 1997 heart attacks killed over 720,000. In general about a third of the heart attack victims do not survive the attack and 44% of the women and 27% of the men who do die within a year. You already know the safeguards against heart attacks. Don't smoke (smoking is an enormous risk factor; it is responsible for about one death out of every six). Eat right. Exercise. Keep your weight and your blood pressure down.

#### Cancer

Cancer is a blanket term that covers a number of different diseases that involve unregulated growth of cells. The probability that an American male will develop cancer at some point is 1/2 and the probability that an American woman will is 1/3. The causes of cancer aren't fully understood, though it is clear that there are different risks factors for different types of cancer, so there are no universal precautions. There are, however, relatively easy ways to decrease your risk of some kinds of cancer (smoking is a very large risk factor for lung cancer and several other cancers; spending long hours in the sun is a risk factor for skin cancer).

#### Strokes

A stroke (brain attack) occurs when the blood supply to the brain is blocked (typically by a clot). About 600,000 Americans suffer strokes each year, and in 1997 almost 160,000 died from strokes. Hypertension (high blood pressure) is an important risk factor for strokes, and in general the basic rules apply here too. Don't smoke, eat sensibly, exercise, treat high blood pressure, keep your weight down.



### Other Causes of Death

Some of the other leading causes of death will probably surprise you. For example, pneumonia and influenza is sixth on the list, diabetes is seventh, and blood poisoning isn't all that far behind.

### Risk Factors

It is always possible to get more informative risk ratios by making the target group more precise. Instead of looking at the rate of strokes in the entire population we could look at the rate of strokes by age group, e.g., from 20-30, 21-40, etc. There is a tradeoff here between more precision and more complicated statistics. But the general idea here is important. About one in three Americans will die of heart disease, but the risk is much higher in some groups than in others. If several members of your family had heart disease, your risk is higher; if you are over fifty or overweight the risk goes up. Again, strokes are the third leading cause of deaths among Americans, but over two thirds of stroke victims are 65 or older.

Just as we can consider smaller subgroups rather than looking at all Americans, we can consider larger groups by looking at a number of countries or even at the entire world. This will often change risk factors, since the risks facing people in poor third-world countries are often quite different from those facing Americans. For example, the leading world-wide cause of death in 1993 was infectious diseases. They killed 16.4 million people world wide, as opposed to heart disease, which killed 9.7 million. But since about 99% of the deaths from infectious diseases occurred in developing countries, they don't show up as risk factors for Americans.

If you are really concerned about a given risk factor, you can usually find statistics that break the risk down by groups, and you can see what the risk is for the group you are in (e.g., males between 18 and 28 years of age). But even a simple break down like the one in Table 21.1 on page 430 gives us a pretty good idea about risks that can lead to death.

### Smoking

Smoking is the single greatest preventable risk factor in America. About 420,000 people die each year from smoking related deaths (this isn't reflected directly in our table, but smoking leads to items, like heart attacks and cancers, that are on the list). Males smokers reduce their life expectancy by over eight and a half years and female smokers reduce theirs by over four and a half year. Being overweight is also a major risk factor for several of the leading killers (not just heart attacks but also cancer).

## 21.4 Crime Risks

In 1996 Oklahoma City was ranked the third most dangerous city in America (behind New Orleans and Albuquerque; it was actually probably fourth, because Miami didn't report statistics for that year).

What are your chances of being the victim of a crime? It depends on many things: how old you are, where you live, what risks you take. It is possible to break the statistics down for each of these categories, but we won't go into that level of detail here. Two general types of statistics are relevant in thinking about crime: crime rate and victimization rate. The figures aren't extremely precise, because many crimes go unreported, but they are in the right ballpark.

Violent crimes have been decreasing over the last few years, but they are still a very real risk in many parts of American. In 1997 there were almost 19 thousand homicides (considerably less than the number of suicides), about 95,000 rapes, and over half a million robberies. But when you think about risks from crime, you will be more interested in victimization rates. The following table (from the *New York Times 1999 Almanac*, p. 309) gives the victimization rates in 1996 for several violent crimes; the figures report the number of victimizations per 1,000 persons.

1. Simple assault: 42
2. Aggravated assault: 26.6
3. Robbery: 8.8
4. Rape and sexual assault: 5.2
5. Household burglary: 205.7
6. Auto theft: 13.5

Table 21.2: Victimization Rates for 1996

But these statistics vary a great deal for different groups. Your chances of being robbed if you work the night shift at the 7-11 are much higher than the national average. Your chances of being murdered if you live in the inner city are much higher than average. The victims of most crimes are most likely to be black, poor, young, and urban.

It is important to be clear about the difference between a *crime rate* and a *victimization rate*. A crime rate tells us what percentage of people *commit* a given type of crime, e.g., how many people commit assault. A victimization rate tells us what percentage of people are victims of a given type of crime, e.g., how many are assaulted.

### 21.4.1 Exercises

1. Express the figures in the table of victimization rates (Table 21.2 on the preceding page) in terms of probabilities.
2. Find the victimization rates for homicide and arson (use the web).

## 21.5 Other Risks

### 21.5.1 Sex

Most of the risks here are well-known and easily avoidable. But lots of people don't manage to avoid them. The two major risks here are various diseases, especially AIDS, and unwanted pregnancies.

### 21.5.2 Love and Marriage

There is a very real risk that if you get married you will end up getting divorced. According to the *Oklahoma Gazette* (Nov 20, 1997) Oklahoma has the *second* highest divorce rate in the nation (trailing only Nevada). A 1987 University of Wisconsin study based on 1987-89 data found that in the country overall, 27% of married couples divorced in the first decade after their marriage. The rate has declined since then, but this is only a ten year period. Of course few people think that they will be among the casualties, but many of them will be. This isn't a reason to stay single, but it reminds us that nothing is without its risks.

### 21.5.3 Jobs and Businesses

You are *much* more likely to be injured working in a meat packing plant or on an oil rig than working in a shoe store or an insurance office, but in the modern Western world, most jobs are relatively safe. There isn't a lot of risk. It's a different story if you are thinking of starting a new business. This isn't to say that you shouldn't start a new business; many prosper and thrive. But one out of five businesses go bankrupt each year, and one out of two new businesses go under within a decade.

## 21.6 Cognitive Biases and the Misperception of Risk

All sorts of things make it easy to misperceive risks. The media report certain types of calamities (e.g., people killed in fires) more often than others that are in fact more common (e.g., drownings). Then too, the more grisly cases stick in our minds. And as if that weren't enough, there are people who have a vested

Many fallacies and cognitive biases lead us to misperceive risks

interest in exaggerating certain risks (You need more insurance; You've got to take this special dietary supplement to avoid liver cancer). Finally, risks that are really serious (like heart disease) may require big changes in our lives, so it is often tempting to downplay them.

Many of the biases and fallacies we have studied lead us to overestimate the risk of some things and to underestimate the risk of others.

### **Sample Size and Bias**

Whenever we draw inferences from small or biased samples, our conclusions will be unreliable. This is as true when the conclusions are about risks and remedies as it is about anything else.

### **Neglecting Base Rates**

If we neglect base rate information our estimates of various outcomes can be highly distorted. Often we hear figures that sound very dramatic, but they sometimes become trivial when we learn about the relevant base rates. A new drug cuts the death rate from the Plague in half. But the base rate of Plague is very low (about four Americans got it last year).

Whenever we hear about risks our first question should always be: *What is the base rate?* Typically we don't need a very precise answer; a ball park figure is usually enough.

### **Availability**

If we don't appreciate how large the difference between the probabilities of having a heart attack and the probability of dying at the hands of a terrorist are, even after September 11, it will be difficult to make rational plans about diet and travel. Several thousand Americans will die from heart disease in a given year, whereas in every year but 2001 about one in a million Americans die at the hands of terrorists.

Things can be available for many different reasons. The media report some things more than others (e.g., fires more than drownings; plane crashes more than car wrecks). Moreover, people you know tend to talk more about some risks than others (if your uncle was recently mugged, you will hear a lot about muggings).

Sometimes a particularly vivid and horrifying sort of accident comes to mind more easily simply because it is more frightening. For example, being electrocuted by wiring in your home sounds very gruesome. But only 200 Americans (less than one in a million) a year die from electrocution. By contrast, over 7000 die from falls in the home. And of course no one can forget the sight of the twin towers at

the World Trade Center collapsing.

### **Probabilities of Conjunctions and Disjunctions**

We tend to overestimate the probabilities of conjunctions and underestimate the probabilities of disjunctions. This means that we underestimate likelihood of failure and overestimate the likelihood of success. This can lead us to underestimate the likelihood of certain risks.

### **Cumulative Effects**

We are prone to underestimate the power of cumulative effects. For example, a contraceptive device may work 99% of the time, but if we rely on it frequently over the years, there is a good chance that it will eventually let us down. Suppose, for example, that you use a particular brand of condom that breaks 1% of the time. Each time you use it the chances are low that it will break. But if you use that brand of condom a couple of hundred times, the chances of a failure start to mount up. Similarly, the chances of being killed in an automobile accident each time we drive are low, but with countless trips over the years the, odds mount up.

### **Coincidence**

Wilbur survives a disease that is fatal to 99.8% of the people who contract it. Wilbur's case is rare, and so people will talk about it, it may make the papers or T.V., and so we are likely to hear about it. If we focus too much on the lucky few who survive a disease despite doing everything their doctor warned them not to, we may conclude that the risk of the disease is much lower than it actually is.

### **Regression to the Mean**

If we overlook regression to the mean, we may think that certain measures will decrease targeted risks, even when they are ineffective and only seem useful because they happened to coincide with regression to the mean. For example, we may overestimate the power of a given policy (like increasing the number of police or enacting tougher sentencing laws) to cut down on crime. This will mean that we have an inaccurate perception of the risks of various crimes and the best ways to combat them.

### **Illusory Correlation**

When we believe in an illusory correlation we think that changes in one thing tend to accompany changes in another. For example, we may think that certain jobs or occupations have a higher (or lower) correlation with various diseases than they really do. This will lead us to overestimate (or underestimate) the risks of various undertakings.

### **Anchoring and Adjustment**

It is possible to set anchors at unreasonably high, or unreasonably low, probabilities for a given type of risk. Even though we frequently adjust for these anchors, we often don't adjust enough. So a high anchor can lead us to overestimate the likelihood of a risk and a low anchor can lead us to underestimate it.

### **Wishful Thinking and Dissonance Reduction**

It is often easier to deal with a risk by convincing ourselves that it's not really as serious as other people say. When the Surgeon General's first report on the dangers of smoking came out in 1964 only 10% of nonsmokers doubted the report. 40% of heavy smokers did.

It's easy to dismiss a report of a recent study suggesting that one of our favorite foods causes cancer by saying that everything causes cancer and the experts keep changing their minds anyway. This is not an unreasonable reaction to a single study. But many of the greatest health risks, e.g., smoking and heart attacks, are established beyond all reasonable doubt. Unfortunately the remedies, while having little financial cost, can exact a huge cost in the changes of lifestyle they require. Many people who would pay a lot of money to avoid these risks won't pay the price of lifestyle change. It is easier to downplay the risk.

### **Framing Effects Revisited**

Earlier we learned that people are typically *risk averse* when it comes to possible gains. We prefer a certain gain (say of \$10) to a 50/50 chance of getting \$20 (even though these alternatives have the same expected value). In fact, many people prefer a certain gain (say of \$10) to a 50/50 chance of getting \$25 or even more.

By contrast, people tend to be *risk seekers* when it comes to losses. Most of us prefer the risk of a large loss to a certain loss that is smaller. For example, most people prefer B (a 25% chance of losing \$200, and a 75% chance of losing nothing) to A (a 100% chance of losing \$50). How we think about risks often depends on

how things are framed. More specifically, it depends on whether they are framed as gains (200 people are saved) or as losses (400 people die).

When we frame a choice in terms of a certain loss we think about it differently than we would if we frame it in terms of insurance. When we frame a choice in terms of people being saved we think about it differently than we would if we frame it in terms of people dying.

### **21.6.1 Tradeoffs**

Often the only way to decrease one risk is to increase another one. To take a whimsical example first, you will decrease your risks of being hit by a car or falling under a train if you stay home all day in bed. But in the process you will have increased the risks of injury from falling out of bed, the risk of countless health problems due to lack of exercise, and the risk of being poor since you'll probably lose your job.

The same point holds for risks that are a serious worry. Many illnesses are best treated with medication, and if they are serious you may be better off in a hospital. But it is well known that hospitals are breeding grounds for infection. It is less well known, but according to an article in the *The Journal of American Medical Association* published in April, 1998, adverse drug reactions (ADRs) in American hospitals may be responsible for more than 100,000 deaths nationwide each year. This would make adverse reactions to prescribed medications a very significant cause of death.

## **21.7 Psychological Influences on Risk Assessment**

### **21.7.1 Individuals Differences**

People differ greatly in their attitudes toward risk. Some people enjoy risks, and other people will go to great lengths to avoid risks. In general people are more willing to put up with voluntary risks than with risks that are imposed on them. For example, many people are willing to take a fairly large risk when they go rock climbing (since they chose to do it), but they would be very upset by the (probably much lower) risk brought on when the government decides to put a toxic waste dump near their town.

There isn't anything irrational about this; it probably reflects important facts about personal autonomy. But people also tend to perceive something as less risky when it is voluntarily incurred, which is simply a misperception.

People are also more tolerant of risks that they have some power to deal with than with risks over which they have no control. Many people feel safer driving a

car (control condition) than riding in the passengers seat (less control). They also perceive such things to be less risky. People also perceive natural risks to be less severe than man-made risks, and they think of risks involving novel technology or especially dreaded outcomes (like nuclear power facilities) as especially great.

### 21.7.2 Groups

Later in the book we will read about the risky shift. The *risky shift* occurs when people who take part in a group discussion are willing to support riskier decisions than they would individually, before the group discussion.<sup>1</sup>

## 21.8 Chapter Exercises

1. Here are answers to some of the pretest at the beginning of the chapter. A + indicates that the cause is the more likely of the pair and a - that it is the less likely. Use the information in the Table of causes of death (Table 21.1 on page 430) to fill in the number of deaths per hundred million people. The figures for the causes not listed in that table are given below.

- |                        |                     |
|------------------------|---------------------|
| 1. 1. Diabetes +       | 2. Homicide -       |
| 2. 1. Suicide +        | 2. Homicide -       |
| 3. 1. Asthma +         | 2. Tornado -        |
| 4. 1. Lightning -      | 2. Flood +          |
| 5. 1. Lung Cancer -    | 2. Heart Disease +  |
| 6. 1. Stomach Cancer + | 2. Rabies -         |
| 7. 1. Homicide -       | 2. Stomach Cancer + |
| 8. 1. Suicide +        | 2. Syphilis -       |

The numbers after each cause of death give the approximate number it kills a year per one hundred million Americans.

1. Asthma: 920
2. Tornado: 44
3. Lightning: 52
4. Flood: 100
5. Lung Cancer: 37,000

---

<sup>1</sup>Larry Laudan has written two easy-to-read books on risk: *The Book of Risks*, John Wiley & Sons, 1994, and *Danger Ahead: The Risks you Really Face on Life's Highway*, John Wiley & Sons, 1997. A somewhat more technical discussion will be found in Peter G. Moore's *The Business of Risk*, Cambridge University Press, 1983.



6. Stomach Cancer: 46,600
7. Rabies: < 1
8. Syphilis: 200

2. How likely would you think it is for someone in the U.S. to die from anthrax in a given year? How likely is it for them to drown? You don't need exact numbers, just good ball park figures. Why are most of us so much more frightened of the use of anthrax by terrorists?

3. First, give your best estimate of each of the following:

1. being killed in a car wreck in a one year period.
2. being killed in a boating accident in a one year period.
3. being killed while riding a bicycle in a one year period.
4. being killed in a plane crash in a one year period.
5. being killed in a fire in a one year period.
6. being murdered in a one year period.

Then go to the web and find out the actual probabilities. How close were you. Try to explain why you were wrong, in any cases where you were way of the mark.



## **Part IX**

# **The Social Dimension**



## Part IX. The Social Dimension

Human beings are social animals, and the thoughts and actions of others have an enormous impact on our own actions and thoughts. In this module we will examine several aspects of the social dimension of reasoning.

In Chapter 22 we examine the most central, non-rational ways in which other people influence our attitudes and thought. We acquire many of our most deeply-rooted attitudes and beliefs in the process of growing up in the family and society that we do. As we mature, the pressures of peer groups, professional persuaders (like advertising agents), and authority figures influence our attitudes and thoughts, often without our even realizing it. Social influences often have good consequences, but as cases like the Holocaust show, they can also lead to terrible results. The goal in this chapter is to become more aware of these social forces and to devise safeguards to diminish their power over us.

In Chapter 23 we turn to our attempts to understand and explain human behavior. We will find that people often greatly underestimate the power of the context or situation in which *other people* act; we attribute their actions to their traits or desires or abilities, when in fact the situation in which they act plays a bigger role in explaining their behavior. We will also examine the differences between the causes people find for their own actions and the causes they find for the actions of others. The goal in this chapter is to do a better job at explaining why people do the things they do.

In today's world many people's jobs require them to work as a part of a group, projects are carried out by teams, and numerous decisions are made by committees. Juries, legislative bodies and, most importantly families, are groups that must think about goals and make group decisions. There is great variability among groups, so we can't expect any simple, blanket conclusions that apply to all of them, but we will see that groups are susceptible to several sorts of biases. In Chapter 24 we learn more about these biases and shortcoming in group thinking and develop ways to avoid them.

The harmful effects of biases and prejudices in our society are all too obvious. There are *many* reasons why people are biased against members of certain groups, but one important cause of prejudices and stereotypes is *faulty reasoning*. The goal of Chapter 25 is to see how the biases and fallacies that we studied in earlier chapters foster (and help maintain) prejudices and stereotypes. It would be too much to hope that clearer thinking would eliminate such problems, but it would be a step in the right direction.

Social Dilemmas are situations in which actions that seem to be in each individual's own self-interest lead to outcomes that are worse for everyone. Such situations occur in international relations, in countries and cities, in families, and in two-person interactions. The goal in Chapter 26 is to learn about the causes of such dilemmas and to examine strategies for extricating ourselves from them.

# Chapter 22

## Social Influences on Thinking

**Overview:** Human beings are social animals. The thoughts and actions of others have an enormous impact on our own actions and thoughts. Social influences often lead to good reasoning, but when we rely on them too much, or in the wrong context, their influence can be disastrous. They can be the worst impediment there is to clear and independent thinking. In this chapter we will examine several ways that other people can influence our actions and thoughts and discuss some remedies for the problems that these pose.

### Contents

---

<b>22.1 The Social World</b> . . . . .	<b>450</b>
<b>22.2 Persuasion: Rational Argument vs. Manipulation</b> . . . . .	<b>450</b>
<b>22.3 Social Influences on Cognition</b> . . . . .	<b>452</b>
<b>22.4 Socialization</b> . . . . .	<b>452</b>
<b>22.5 The Mere Presence of Others</b> . . . . .	<b>453</b>
<b>22.6 Professional Persuaders</b> . . . . .	<b>454</b>
22.6.1 Professional Persuaders: Tricks of the Trade . . . . .	454
<b>22.7 Conformity</b> . . . . .	<b>457</b>
22.7.1 The Autokinetic Effect . . . . .	458
22.7.2 Ash's Conformity Studies . . . . .	458
<b>22.8 Obedience</b> . . . . .	<b>461</b>
22.8.1 The Milgram Experiments . . . . .	461
22.8.2 Changing Behavior vs. Changing Beliefs . . . . .	464
<b>22.9 What could Explain Such Behavior?</b> . . . . .	<b>464</b>

---

22.9.1 Obedience Training . . . . .	464
<b>22.10 Responsibility . . . . .</b>	<b>465</b>
22.10.1 I Was Just Following Orders . . . . .	465
22.10.2 Asleep at the Wheel . . . . .	465
<b>22.11 Safeguards . . . . .</b>	<b>465</b>
22.11.1 The Open Society and the Importance of Dissent . . . . .	466
<b>22.12 Chapter Exercises . . . . .</b>	<b>467</b>

---

## 22.1 The Social World

We can't help but rely on other people. Most of the knowledge we have was acquired from others. Most of the goods we own were made by others. In a diverse, highly technological society like ours, we continually have to rely on the expert opinions of others. And the most important things in life for most of us are our relationships with other people.

Given the importance of the social world for almost all aspects of our lives, it is not surprising that social factors influence our thinking in profound ways. These influences often promote good critical reasoning. But they also have their dark sides, often leading to poor reasoning; indeed, as Nazi Germany shows, they can lead to beliefs and actions with terrible consequences.

It isn't possible to completely escape such influences, but understanding how they work can help us guard against their undue influence. Moreover, to the extent that we are aware of them, we are less vulnerable to those who are willing to use this knowledge for their own ends.

## 22.2 Persuasion: Rational Argument vs. Manipulation

We often find ourselves trying to persuade or convince other people of one thing or another. People in some jobs do this for a living, but no matter what our vocation, we are likely to do it. You might want to convince your teacher that you deserve a second chance on the big exam, or to convince your students that they should care about critical reasoning. You might want to convince someone to go out on a date with you, or to marry you, or to divorce you. You will almost certainly want to convince your children, once they are old enough to understand, that hurting other people for no good reason is a bad thing to do. In fact, we spend a lot of time and energy trying to convince other people.



There are many different (and often subtle) techniques for persuading people of things. One of the main points in this course is that the best way to do this is by giving them a good argument that employs premises they accept. Why does this matter? If we do this, we treat other people as autonomous adults who we think can make up their own minds. We give them what we think are good reasons, and then let them make up their own minds. But if we try to persuade them in other, nonrational ways, we treat them as objects to be manipulated (I'll say whatever he wants to hear, if it will get him to buy this car) or as children who aren't capable of thinking for themselves ("After all, I know what's best for him").

The latter approach is called *paternalism*. It assumes that other people are not capable of thinking for themselves. This is a very sensible view to take about young children, and we often have to extend it to adults who suffer from severe mental disturbances or who act in ways that harm others (though there is much debate about just who falls into this category). But it's a very dangerous view to take about normal adults in general.

Just a little thought about the history of the twentieth century should convince us of the dangers of deciding that others don't know how to reason correctly or how to decide what is best for them. An easy way to see why this is objectionable is to think about how we would feel if other people treated us as an object to be manipulated or a child to be cajoled and tricked into acting and thinking in the ways others want us to.

Life is too short for us to devote hours thinking about each decision we make. But when the decisions are important, we should think about them for ourselves. Even in these cases, rationality is an *ideal*. In this respect it is like a good marriage: it's a goal well worth striving for, even though there will be lots of lapses and backsliding and even on our best days we won't fully achieve it.

Of course there are many other ways to persuade people. Indeed, we have encountered a variety of techniques that can be quite effective for doing this. One of the most effective ways to do so is to provide what *seems* like a good argument on the surface, but which persuades (if it does) because it takes advantage of various cognitive biases (e.g. our tendency to ignore base rates) or because it appeals to our emotions or self-interest. This is one reason why the study of fallacies and cognitive biases is worthwhile. Our earlier modules cover many of the ways in which bad arguments can persuade us when we aren't careful. We will now turn more directly to the social aspects of persuasion.

### 22.3 Social Influences on Cognition

There are many ways in which other people influence our own behavior and attitudes. Many of them involve one or more of the following:

**Socialization** We acquire many of our most fundamental beliefs, attitudes, and values in the very process of growing up.

**Experts** We constantly rely on the views and advice of experts, including teachers, text books, and much of the mass media. We have examined the role of experts in detail in an earlier chapter, so we won't discuss them here.

**Mere Presence** Our performance on cognitive tasks is affected by the mere presence of others; even as a passive audience can influence how well we do.

**Persuasion Professions** We are targets of people in the "persuasion professions" (advertising agents, lobbyists, politicians, social reformers), who are constantly trying to influence how we think.

**Conformity** We are very strongly influenced by the views and actions of our peers (or those of members of groups we admire)

**Obedience** We are all more susceptible to the views and commands of those in authority than we would suppose.

Most of these influences can be useful, and none of them are intrinsically bad. We often have to rely on the views of experts; our society wouldn't do well without rules and authorities with the power to enforce them (though perhaps some of the specific rules are flawed). Sometimes these things only influence our actions, but often they influence our thoughts, attitudes and beliefs (frequently without our even realizing it). In the rest of this chapter we will examine some of the ways in which these influences operate, and we will devise some safeguards to diminish their power over us.

### 22.4 Socialization

Most people's basic picture of the world is largely determined by what they were taught as they grew up. A young and helpless infant doesn't have the tools to question the things her parents teach her; until she has mastered a certain amount of language, she doesn't possess the words or concepts required to frame challenges or doubts. As we acquire language we get a set of categories and principles for

thinking about the world. As we are rewarded and chastised we acquire a sense of what is right and what is wrong.

In our early years we absorb many of the beliefs of the people raising us, entirely unaware we are doing so. As we grow older, additional social forces come into play: teachers, peers, the mass media, and so on.

Think of the beliefs that are most important to you. They are likely to include beliefs about things like religion, morality, patriotism, love. Can you recall a time when someone reasoned with you and got you to change your mind about these matters? Did you ever seriously entertain the thought that some religious or moral views quite different from your own might be true and that yours might be false? If you had grown up in a culture with a very different religion or morality, what do you suppose you would now think about these matters?

Certainly people sometimes change their views about such matters; some are converted to religion or come to see it as much more important than they had (they are “born again”), while others lose their faith. But many people acquire such beliefs when they are very young, and they retain them with little alteration for the rest of their lives. The fact that you acquired a belief simply because your parents held it doesn’t mean that it’s false.

But *if* you continue to hold it solely because other people taught you to do so, you are handing control of your mind over to others. Of course no one has time to constantly examine all of their fundamental beliefs. But it is healthy to examine some of them now and again, and the years you spend at a university are a particularly good time to do it. You may finish the process with exactly the same views you have now. But if you have thought about them critically, they will then be *your* views.

### Exercises

1. Can you remember a time when you didn’t hold pretty much the beliefs that you have now about religion? About morality?
2. If you had grown up in a very different culture, one with a quite different religion or morality, what do you suppose you would now think about these matters?

## 22.5 The Mere Presence of Others

The mere presence of others can affect our performance on many sorts of tasks. **Social facilitation** occurs when their presence enhances performance. For example, many people do better in athletic events if an audience is present. The same holds

true for cognitive tasks; we often do a better job at solving verbal or mathematical problems and puzzles if others are watching.

In some cases, however, the presence of others detracts from our performance. This is known as **social impairment**. Research suggests that an audience enhances someone's performance on a task if they are accomplished at it, but it detracts from their performance if they are not. Although these findings are of interest, we will focus more on longer-term social influences on thought.

## 22.6 Professional Persuaders

Many people work in the **persuasion professions**. The success of professional advocates like advertising agents, lobbyists, spin doctors, trial lawyers, and politicians as well as the success of social reformers and charity workers depends on their ability to persuade others to do something. Their goal is to convince us to buy Wheaties or to vote for George W. Bush or to contribute to the March of Dimes.

Often people in the persuasion professions have a bad reputation: the stereotypical used-car salesman would run over his own mother to clinch a deal. But in many cases professional persuaders are admirable: the world is a better place because of those who try to convince us to give some of our time or money to those in need.

Often the goal of professional persuaders *is* to manipulate our beliefs or attitudes in a way that will benefit *them*. For example, political life is increasingly a matter of advertising and image manipulation. Nowadays many candidates are marketed like corn flakes, their message fine-tuned to reflect the latest poll results, their every word explained by spin doctors.

### 22.6.1 Professional Persuaders: Tricks of the Trade

There are many techniques for persuading people. Some involve pressuring them, but the most effective devices are the ones that people don't even notice. For example, an real-estate agent might exploit the contrast effect by first showing a prospective buyer a run-down over-priced house right before showing the house he is really trying to sell. In this section we will learn about three of the most effective techniques for getting people to do things without their even realizing that they are being manipulated.

#### The Foot-in-the-Door Technique

One very effective device is the **foot-in-the-door technique**. The foot-in-the-door technique involves getting someone to do or believe something that is reasonably

small. After they do agree to the small request, the person is more likely to comply with a larger request or suggestion. Professional fund raisers are well aware of this technique. Often they first ask for a small donation, then come back later to ask for a larger one.

In a real-life study a number of housewives were asked a few questions about which soaps they used. A few days later both the group who had answered these questions and another group who had not been contacted before were asked if a survey team could come to their home and spend two hours recording every product that they owned. Housewives who had agreed to the small requests (to answer a few questions about soap) were over twice as likely to accede to the much larger request.

**Foot-in-the-door technique:** first get someone agree to something small

An even more dramatic illustration of the foot-in-the-door technique comes from a 1966 study a group of residents of Palo Alto. Psychologists going door-to-door asked a number of residents to display a modest three-inch sign saying BE A SAFE DRIVER. Two weeks later another person was sent around, both to the people contacted earlier and to another group of people who hadn't been. He asked for permission to erect an enormous billboard on the resident's front lawn that proclaimed DRIVE CAREFULLY and showed them a picture clearly depicting the billboard as an enormous monstrosity. Only 17% of the people who had not been contacted before agreed to the request. But 55% who had been contacted earlier and displayed the small, three-inch sign agreed. In other words, over half of those who had acceded to the earlier, smaller request, agreed to the bigger one.

The foot-in-the-door technique is commonly used by people in the persuasion professions. A salesman at the door often asks for something small like a glass of water. Once the resident agrees to that request, the salesman has a better chance of getting them to buy something. There aren't many door-to-door salesmen nowadays, but telemarketers have adopted this technique too. It was also used, less innocently, by the Chinese during the Korean War. They made small, innocent sounding requests of their prisoners of war, and then moved very gradually on to larger requests.

### Lowballing

There is a related phenomenon known as **lowballing**. This occurs when a person is asked to agree to something on basis of incomplete or inaccurate information about its costs. Later she learns that the true cost is higher. But having made the original commitment, she is more likely to accept the new cost than she would have been had she known about it up front.

**Lowballing:** getting a person to agree to something on the basis of inaccurate information about its costs

For example, car dealers sometimes clinch a sale, go off to verify it with their boss, then return with the news that it's going to cost just a little more than they'd

thought. Someone might ask you for a ride home and when they get in your car announce that they live twenty miles away. In both cases, the person who made the original commitment is more likely to follow through on it than they would have been had they known about its cost at the time that they made their decision. The subjects made an initial commitment to be in the experiment and only later discovered what they had gotten themselves into.

### The Door-in-the-Face Technique

**Door-in-the-face technique:** first get someone to refuse a large request

The **Door-in-the-face technique** is another device for eliciting compliance. The strategy here is to lead someone to believe or do something by first asking them to do something bigger (or to believe something less probable), which you know they will refuse. After the larger request is refused, the person is often more likely to do or believe the second, smaller, thing.

Robert Cialdini and his coworkers asked one group of people to volunteer to work as a counselor for two hours a week in a juvenile center for at least two years. Not surprisingly, no one agreed to this. Later the people in this group and an equal number of people who had not been contacted before were asked to chaperone a group of juvenile delinquents on a trip to the zoo. People who had first been asked the much larger request (to become a counselor for two years) were over three times as likely to take the delinquents to the zoo as those who had not been asked.

In another study subjects were asked to contribute time to a good cause. Some of them were asked to contribute a lot of time. Most refused, but they were later asked to commit less time. Only 17% of those who were only asked for a small amount agreed, but 50% of those who were first asked for a large amount agreed to a smaller amount. This technique is common in bargaining and negotiating at all levels, from negotiations between nations to negotiations between parents and children.

### Safeguards

People can resist these pressures, but it requires some thought. We often go along because we act without paying much attention to things. The following study illustrates the point. There was a line to use the only nearby photocopying machine at a library. A group of psychologists had a person ask to cut in front of others in the line. If the person simply asked to cut in without giving any reason, most of the people in the line refused. And if they gave a good reason like "Could I cut in because I'm running late to pick up my child from school?" most people let them cut in. No surprises thus far.

What is surprising is that if the people conducting the study gave anything that

had the format of a reason, most people let them cut in. For example, if they asked “Could I cut in because I need to make some copies?” many people granted their request. These people’s minds weren’t really in gear, and the mere fact that it sounded vaguely like a reason, though it wouldn’t have seemed a good reason if they had thought about it, was all it took for them to allow someone to cut in front of them.

## 22.7 Conformity

We comply with someone’s wishes if we do what they ask us to do. Conformity involves a more subtle pressure, and we are often unaware of its influence. Conformity may result either from a desire to be right (this is sometimes called **informational influence**) or from a desire to be liked or to belong or to seem normal (**normative influence**).

Informational influence leads to what is sometimes called *social proof*. We use social proof when we attempt to determine what is correct by seeing what other people think is correct. This is often a good way to proceed. If we aren’t sure which fork to use for the salad at a fancy party, it is natural to see which fork others use. If we aren’t sure what the speed limit is, it makes sense to match our speed to the average speed of other drivers. And in ambiguous situations we often think that other people have a better idea of what is going on than we do (while they may be thinking the same thing about us).

**Social proof:** using what others think to try to determine what we should think

But social proof isn’t always good. When people are about to vote on some important issue in a meeting some of them first look around to see how others are voting and then try to go with the majority. And under the wrong conditions social proof can lead to disaster. One of the reasons many Germans went along with the Nazis was that many other Germans did so too.

In the case of normative influence we go along to get along, to be liked or accepted or at least not despised. Normative influence can lead people to conform publicly, but they may not privately accept the views they act like they have accepted. Sometimes normative influence only involves an isolated action, but it can involve norms that affect us on many occasions.

**Norms** are explicit or implicit rules that tell us what sorts of behavior, attitudes, beliefs and even emotions are appropriate in a given situation. For example, in our society it is appropriate to be angry under some circumstances (e.g., when we see an injustice committed) but not others (e.g., if someone unintentionally mispronounces our name).



### 22.7.1 The Autokinetic Effect

Muzafer Sherif (1906-1988) was one of the pioneers of experimental social psychology. He was born in Turkey, but came to the United States and did much of his work at the University of Oklahoma. One of his most famous studies involved a perceptual illusion known as the autokinetic effect, but the experiment was really about group norms and conformity.

You are blindfolded and led into a dark room; you aren't sure where the walls are or what the size of the room is. If someone shines a tiny spot of light on a fixed spot on the wall in front of you it will appear to move, even though it is completely stationary. This is the *autokinetic effect*; it is a standard perceptual illusion. Of course if people don't know about this illusion, they will think the spot of light really does move and then disappear. But people differ a good deal in *how much* they think it moves. Some think it's just a few inches; others think it's several feet.

Sherif told his subjects that they were participating in an experiment on perception and that their task was to estimate how far the light moved on each of a number of trials. When subjects performed the task by themselves, each developed a characteristic response (it moved two inches; it moved a foot and a half).

In another condition subjects worked in groups of two or three. In these conditions the subjects' estimates of the distance the light moved would converge until they were in very good agreement. *Group norms* emerged. Different groups would settle in on different norms, but the norms within each group were quite stable. In another condition Sherif introduced a confederate into some of the groups who was forceful enough to get the rest of the members to adopt his norms.

The norms established by a person's group persisted for at least a year (when the subjects were brought in and retested), and as members gradually left the group and were replaced by new members, the norms were passed down to later generations. Although the group members did not realize it, there were powerful pressures in the situation that led them to conform.

### 22.7.2 Ash's Conformity Studies

Solomon Ash thought that Sherif had exaggerated the degree to which people conform and in a famous series of studies conducted in the early 1950s he set out to show the limits to conformity. To his surprise, he found something very different.

In a typical Ash experiment a group of experimental subjects were seated in a semicircle around a table. All but one of the people were confederates, accomplices who were in on the experiment. But the lone subject was led to believe that these other people were subjects too.



The people around the table were shown a series of cards, two at a time. The card on the left had a single vertical line, the standard. The card on the right had three vertical lines; one line was the same height as the standard line (the one on the first card), and both of the other lines clearly were not (as in Figure 22.1). The difference was completely obvious to anyone with normal vision, and when subjects in a control condition were shown the cards, only 5% of them made a mistake about which line on the second card matched the line on the first.




Figure 22.1: Which Line Matches A?

The experimenter then asked the people in the group to say which line on the second card was the same height as the line on the first. One by one they gave their verdicts, working their way around the table to the lone subject. Sometimes the confederates gave the correct answer, but sometimes they didn't. When they all agreed in the wrong response, there was pressure on the subject to conform. On 37% of the trials, subjects went along with an (obviously) incorrect response, with about three-quarters of the subjects going along at least once.

It would not be surprising if the rate of conformity differed from one culture to another, and this has been found to be the case. In countries that stress the importance of the group, there is more conformity. A striking thing about Ash's studies is that he found so much conformity in the United States, where the value of individualism is so strongly stressed. Many similar experiments have been conducted in this country over the years. Most find slightly lower rates of conformity, perhaps because many people have become more willing to challenge authority, but Ash's basic results have held up.

Ash found that if there was only one confederate the subject wouldn't conform. But perhaps the most important finding was that if *even one* of the confederates gave the correct response, the subject almost always gave the correct response too. The presence of even one dissenter among a group of conformists was usually enough to undermine the group's influence.



Figure 22.2: Socialization and Conformity

Pressures to conform can be nearly irresistible, and often we go along without giving it any thought, by habit. The dark side to Ash's studies is that *even* when the correct answer was very clear and the subject didn't know any of the people in the group around the table, conformity was common. What would happen if the issue was murky or if the group included one's friends and people she admired?

### Exercises

1. What would you have predicted would have happened in Ash's study if you heard it described but weren't told the results?

2. What would you have predicted that *you* would have done in Ash's study if you heard it described but weren't told the results?
3. What do you think that you would have done if you had been in Ash's study?
4. Why do you think the presence of a single dissenter could radically decrease the amount of conformity?

## 22.8 Obedience

### 22.8.1 The Milgram Experiments

**Obedience to authority** is a particular kind of compliance. It occurs when someone can punish us for disobeying, but it also occurs when the person making the request seems to have a *legitimate right* to do so. The notions of conformity, compliance, and obedience shade off into each other, and we won't be concerned with drawing precise boundaries among them. The important point is that there are many clear cases of each, and in all of them the actions of others can influence us.

We are trained to obey. But not all obedience involves a blind, cringing willingness to do whatever we are told; there are often good reasons for obeying. It would be impossible to raise children if they never did what they were told. And in many cases there are good reasons for adults to comply with the directives of legitimate authorities. If a traffic policeman directs our car down another street, it is usually reasonable to do as he asks.

But sometimes people fulfill requests or follow orders that harm others or that violate their own views about right and wrong. *Excessive deference to authority* occurs when someone gives in too much to authority and quits thinking for themselves. This is what happened in one of the most famous experiments ever conducted.

These experiments were conducted by the psychologist Stanley Milgram and his coworkers at Yale University from 1960 to 1963. Each session two people enter the waiting room of the psychology laboratory. One is a subject. The second person also says that he is a subject, but he is really a confederate (an accomplice who works with the experimenter). The experimenter tells the pair that they are going to participate in an experiment on learning. One of them will play the role of the teacher and the other will play the role of the learner. The two are asked to draw



straws to determine who will get which role, but in fact the drawing is rigged so that the real subject is always the teacher and the confederate is always the learner.

The learners are supposed to learn certain word pairs (e.g., boy-sky, fat-neck), which they are then expected to repeat in an oral test given by the teacher. The teacher is asked to administer an electric shock to the learner every time the learner makes a wrong response (with no response in a brief period counted as a wrong response). The voltage of the shock increases each time in intervals of 15 volts, with the shocks ranging from 15 volts to 450 volts. The shock generator has labels over the various switches ranging from “Slight Shock” up to “Extreme Intensity Shock,” “Danger: Severe Shock,” and “XXX” (Table 22.1). The learners are strapped down so that they cannot get up or move so as to avoid the shocks. They are at the mercy of the teacher.

<b>Volts</b>	15-60	75-120	135-180	195-240
<b>Label</b>	Slight Shock	Moderate Shock	Strong Shock	Very Strong Shock
	255-300	315-360	375-420	435-450
	Intense Shock	Extreme Intensity Shock	Danger: Severe Shock	XXX

Table 22.1: Labels on the Shock Generator

The teacher and learner are put in different rooms so that they cannot see one another, and they communicate over an intercom. After the first few shocks (at 150 volts), the learner protests and cries out in pain. After a few more shocks the learner protests that he has a weak heart. And eventually the learner quits responding entirely. But since no response is counted as a wrong response, the subject is expected to continue administering the shocks.

The teachers typically show signs of nervousness, but when they are on the verge of quitting the experimenter orders them to continue. The experimenter in fact has a fixed set of responses; he begins with the first, then continues down the list until the subject is induced to continue. The list is:


1. Please Continue (or Please go on)
2. The experiment requires that you continue
3. It is absolutely essential that you continue
4. You have no other choice, you *must* go on

If the subject still refused after getting all four responses, he or she was excused from the rest of the teaching phase of the experiment.

How many people do you think would continue giving shocks all the way up to

450 volts? The experimental setup was described to a large number of people, including students, laypersons, psychologists, and psychiatrists. All of them thought that most of the subjects would defy the experimenter and abandon the experiment when the learner first asked to be released (at 150 volts). The experts predicted that fewer than 4 percent would go to 300 volts and that only one in a thousand would go all the way to 450 volts. Indeed, Milgram himself thought that relatively few people would go very far. But the results were very different. In this condition of the experiment, 65% of the people went all the way, administering shocks of 450 volts, and *all* of them went to at least 300 volts before breaking off. Furthermore, these rates were about the same for men and for women (Table 22.2).

Milgram and his colleagues varied the conditions of the experiment, but in all conditions they found much higher rates of obedience than anyone would ever have predicted. For example, nearly as many women subjects continued to the end, where they administered 450 volt shocks to the subject. And when the experiment was conducted in a seedy room in downtown Bridgeport Connecticut instead of the prestigious environment of Yale University obedience was nearly as great.



The more directly involved the subject was with the victim, the less the compliance. And, in a result that fits nicely with Ash's data, if several subjects were involved and one refused to obey, the others found it easier to refuse. On the other hand, many people went along with an authority, even though they didn't understand what was going on, and even when the experiment was done in a seedy office in downtown Bridgeport rather than on the Yale campus. In short, many of Milgram's subjects went all the way, delivering what they thought was a very painful shock to someone who seemed to be in great discomfort and to be suffering from heart trouble. A compilation of the results of several studies is given in Table 22.2; the number of subjects in each condition is 40.

In some conditions 65% of the subjects went all the way

Condition	Men	Women	Bridgeport
Average voltage	405	370	325
Percent who continued to end	65%	65%	47.5%

Table 22.2: Results of Milgram Experiments

In another condition a team of people (all but one of whom were confederates)

had to work together to administer the shocks. One read the word pairs, another pulled the switch, and so on. In this condition if one of the confederates defied the experimenter and refused to continue, only ten percent of the subjects went all the way to 450 volts. As in the Ash studies, one dissenter—one person who refused to go along—made it much easier for other people to refuse as well.

Ethical standards no longer allow studies like Milgram's, but in the years after his work over a hundred other studies on obedience were run. They were conducted in various countries, and involved numerous variations. Most of them supported Milgram's results

In an even more real-life context, a doctor called a hospital and asked that an obviously incorrect prescription be administered to a patient. Although ten out of twelve nurses said they wouldn't dispense such a prescription, twenty one of the twenty two nurses that were called did comply.

## 22.8.2 Changing Behavior vs. Changing Beliefs

In some cases we do things that we don't think we should in order to win the approval of others. For example, many of the subjects in Ash's conformity studies knew that they were giving a wrong answer, and only did so because of the social pressure they felt. But in many cases we can only get someone to behave in a given way (e.g., to shock unwilling victims) if we change the way they think.

## 22.9 What could Explain Such Behavior?

### 22.9.1 Obedience Training

Children are taught to obey. This is unavoidable. An eight month-old infant doesn't have the language or concepts to understand our reasons for forbidding her to do certain things like pulling the dog's tail. As she matures we can give some explanations ("How would you like it if someone did that to you?"), but it isn't possible to engage in a subtle and detailed argument with a four year old. Someone at this age can of course generate an endless series of "why questions," but at some point the exasperated reply will be "Because I said so."

If people, with their diverse beliefs and goals, are to live together in anything approaching harmony, society must have certain rules, and we are also trained to obey many of these. In some cases the need for rules is dramatic; in war time, for example, it is necessary for people to work together in a coordinated way, and this requires authorities who others will obey. But society also requires such things as laws and taxes to function smoothly. Of course many parents, and many societies, greatly overdo things, but the basic point here is that authority has its place, and

we have been trained to recognize that. The key is for us to think about *whether* we should obey an authority in a given situation.

## 22.10 Responsibility

### 22.10.1 I Was Just Following Orders

Many subjects in the Milgram experiment felt great discomfort as they delivered what they thought increasingly powerful shocks to the learner. They would stop and convey their misgiving to the experimenter. One of the most effective ways of getting them to continue was to assure them that they were *not responsible* for the learner or his health.

When people don't feel responsible for their actions they are capable of doing quite terrible things. After World War II many Germans who had worked at the death camps defended themselves with the refrain "I was just following orders." They were just soldiers doing what soldiers were being ordered to do; the responsibility for their actions lay elsewhere.

### 22.10.2 Asleep at the Wheel

Many of our actions result from habit. Often this is good; we can't stop to think about everything that we do. But the habits to conform and obey make it easy to go along with things we shouldn't, *without really thinking about them*. This is illustrated in a small way in the study where people asked to cut into a line to use a copying machine; it is illustrated in a much more frightening way in the Milgram experiments. One reason subjects went as far as they did in these experiments was because of a habit to obey those who seem to be legitimate authorities. We often do this almost automatically, without thinking about what we are doing. In fact, habit is one of the greatest enemies of clear and critical thinking.

## 22.11 Safeguards

It should go without saying that there are no foolproof ways to avoid the disastrous consequences of conformity and obedience. These have always been with us, and probably always will be. But it doesn't follow that we should just shrug our shoulders and say "That's life." If we can find ways to diminish these catastrophic consequences, that would be very good. There are many things that might help. Here we will consider several that involve reasoning and inquiry.

### 22.11.1 The Open Society and the Importance of Dissent

In the Ash studies, the study involving nurses and improper prescriptions, and Milgram's studies of obedience, the most effective way of reducing conformity or obedience was to have at least one other person present who refused to go along. Often even one dissenter was enough to eliminate conformity or mindless obedience.

This strongly suggests the importance of fostering an atmosphere in which dissent is possible. In a group, including an entire society, where open discussion is allowed, a variety of viewpoints can be aired, abuses by authorities can be exposed, and reasons for resisting them can receive a hearing.

A free and open society, one open to ideas and disagreements, makes critical reasoning much easier. In a society where open discussion is allowed, a variety of viewpoints can be aired. Without free expression the scope of our thoughts will be limited; we will be exposed to fewer novel ideas, and our sense of the range of possibilities will be constricted.

Since no one has cornered the market on truth, we should beware of those who would set themselves up as censors to decide what the rest of us can say and hear. But one price of free and open discussion is having to hear things we may not like. This can be unpleasant, but it can still be a good thing. A view or position may be (1) true, (2) false, or (3) some mixture of the two. In each case we will be better off if a view we find offensive is allowed a hearing.

**1. The view is true** If the view I dislike is true, it should be allowed a hearing. The truth doesn't necessarily set us free, but it does put us in a better position to solve the problems that beset us. Actions and policies based on mistaken views are much less likely to succeed than those based on true views.

**2. The view is false** What if the view I find offensive is false? It might even seem too dangerous for the masses to hear about it—they might be taken in, or led to do things they shouldn't. In addition to the problem of who is to decide which views should be banned (there will always be volunteers for this job), this assumes that most people are so bad at reasoning that they can't be trusted to think about things for ourselves. The claim that others just aren't smart enough to be exposed to certain ideas is both insulting (to them) and arrogant.

Of course some restrictions on free expression are well meaning. Current German laws against denying that the Holocaust occurred and some of the more severe speech codes on today's campuses were instituted with good intentions. But placing certain speech off limits has the effect of suggesting that the views it expresses really *cannot* be refuted. It suggests that our own beliefs are too weak or ill-founded for us to answer our opponents, so that the only option is to ban



expressions of their views.

It is better to allow such speech and then strive to show what is wrong with it. The power of reason in matters of public policy is, unfortunately, limited, but the more our beliefs are based on good reasons the better, and the alternative (letting others determine what we can hear) is worse.

There is a second reason why it is valuable to think about views we don't like. One of the best ways to truly understand our own beliefs is to see how they compare to the alternatives. In trying to meet the challenge of an alternative view, we have to think seriously about what our own beliefs really *mean* and *why* we hold them. This is healthy, because it is very easy for us to hold beliefs that we don't really understand, mouthing slogans and repeating formulas without much comprehension. And if, after careful consideration, we can't give good reasons why our views are better than the alternative, it might be time to modify them.

**3. A mixture of truth and falsity** Complex views or positions usually contain some mixture of truth and falsity. In such cases we can learn something from the part that is right, and we can strengthen our own views by seeing why parts of it are in error.<sup>1</sup>

## 22.12 Chapter Exercises

1. Explain why you think the subjects in the Milgram experiments on obedience did things they would not have done on their own. To what extent did the experimenter simply get them to *do* things (even though they knew they shouldn't) and to what extent did he actually change the way that they *thought* or *reasoned* about things?
2. Could low-balling have played a roles in the Milgram studies?
3. What lessons can be learned from the Milgram experiments? What could be done to make people less likely to do what the subjects in these experiments did? Defend your answers in a brief paragraph.

---

<sup>1</sup>The reasons given here why free speech is important were set forth eloquently by John Stuart Mill about a century and a half ago in his *On Liberty*. It is still a very readable book and is available in several inexpensive paperback editions. Other references: The billboard study was conducted by Jonathan Freedman and Scott Fraser, "Compliance without Pressure: The Foot-in-the-Door Technique," *Journal of Personality and Social Psychology* 4, 1966; 195-203. Robert Cialdini book *Influence: The Psychology of Persuasion*, Revised edition. New York: Morrow, 1993 contains excellent discussions of many subtle techniques for influencing others, including the foot-in-the-door and the door-in-the-face techniques A superb account of Stanley Milgram's studies on obedience will be found in his book *Obedience*, New York: Harper and Row, 1974. For an accessible account of Ash's work see his "Opinions and Social Pressure," *Scientific American*, 1955; 31-35. Further references will be supplied.

4. In what ways might learning about the Milgram experiments change how you think about things? In what ways might learning about the experiments change how you act?
5. If you ended up in a position where you administered shocks to a series of “learners” day after day, what sort of attitude do you think you would develop towards them? How would you think about your own actions and the way that they reflected on you?
6. What light do these experiments shed on the following cases?
  1. The massacre at My Lai, in which American soldiers killed unarmed Vietnamese old men, women, and children, not on the grounds that this was militarily necessary, but because they were ordered to do so.
  2. The Nazi holocaust, and the Nazi trials after the war, where Nazi officers gave as their excuse—or justification—that they were just following orders.
7. What implications does the behavior exhibited in Milgram’s experiments have for issues involving clear, independent reasoning?
8. What things might help one become the sort of person who would resist the experimenter’s orders? (This is a question for each of us, but also a question about how you would want to raise your children)?
9. In Milgram’s study subjects would often continue when told “the experiment requires that you continue,” or “you have no choice.” Why do you think this was effective? How do you think the subjects thought about such remarks?
10. When subjects in the Milgram experiment began to have doubts, the experimenter said that he would “take the responsibility.” What role did this play in their actions? How do issues about responsibility affect our own actions?
11. Canned laughter is commonly used on TV sitcoms and other shows. Experiments have found that when the material is even a little funny, the use of a laugh track leads an audience to laugh longer and more often and to rate the material as funnier. Why could explain why people—most of whom profess to dislike canned laughter—laugh more when there is a laugh track? Do you think that they really perceive the material funnier?
12. Bartenders often “salt” their tip jars with a few dollar bills at the beginning of their shift to simulate tips left by previous customers. Why do you think they do this? Would you be more likely to leave a tip if they had? Give some other examples of this sort of phenomenon.

13. Wilbur and Wilma, both seventeen years old, are trying to negotiate a curfew with their parents.

*Wilbur's scenario:* Wilbur asks his dad if he can stay out until 10:00 on a Tuesday night to go to a basketball game. His dad agrees. Wilbur comes home in time, and later that week he asks his dad if he can stay out until 1:00 on Saturday night to go to a concert. His dad agrees.

*Wilma's scenario:* Wilma asks her mom if she can stay out until 2:00 (am) on a Tuesday night to go to a basketball game. Her mom says, "No way!" Later that week, she asks her mom if she can stay out until 1:00 on Saturday night to go to a concert. Mom agrees.

What is going on in each scenario (the answers may be different in the two cases)?

14. In obedience and conformity experiments, it has been found repeatedly that the presence of even one dissenter (a person who refuses to comply) makes it easier for others to refuse compliance as well. Why is that? Defend your answer. You may find it useful to relate your answer to one or more of the experiments discussed in this chapter.

15. The habitual use of the device of social proof is not conducive to good reasoning. But can you think of some ways in which it might be used to promote good or healthy behavior? Explain your answer.

16. Do you think that the subjects who played the role of the teacher were unusually sadistic or somehow worse than people in general? If not, why did they do what they did?

17. Give an example, either imaginary or from your own experience, of the door-in-the-face technique. Then give an example of the foot-in-the-door technique. In each case, explain why you think the technique worked (if it did) or failed to work (if it failed).

18. Give an example, imaginary or real (it could be from your own life, but you don't need to attribute it to yourself) where a person's ways of thinking and reasoning (as opposed to just their behavior) are changed in an effort to conform.

19. Give an example, imaginary or real (it could be from your own life, but you don't need to attribute it to yourself) where a person has inconsistent beliefs and engages in dissonance reduction to try to eliminate the dissonance or tension generated by the inconsistency.

20. Give three examples of legitimate authorities and cases where it would be reasonable to comply with their requests. Defend your answers in a brief paragraph.

21. Give an example where there would be pressure to comply with the request of an authority, but where it would not be a good thing to do so. Defend your answer in a brief paragraph.



(a) Mass Suicide at Jonestown, 1978

(b) Rescue Workers after 9/11

Figure 22.3: Social Influences: The Bad and the Good

## Chapter 23

# The Power of the Situation

**Overview:** We have seen repeatedly that context strongly influences reasoning. Since we spend much of our time in *social* contexts, it is not surprising that many features of social situations exert a strong influence on our thought and behavior. In this chapter we will see how most of us frequently underestimate the power of situations—typically social situations—to influence our actions and thoughts. We will also learn about some common biases in our reasoning about people, including ourselves.

### Contents

---

<b>23.1 Case Studies</b> . . . . .	<b>471</b>
<b>23.2 The Fundamental Attribution Error</b> . . . . .	<b>477</b>
23.2.1 Explaining why People do what they Do . . . . .	477
<b>23.3 Actor-Observer Differences</b> . . . . .	<b>482</b>
<b>23.4 Special Cases</b> . . . . .	<b>482</b>
<b>23.5 Chapter Exercises</b> . . . . .	<b>484</b>

---

### 23.1 Case Studies

#### Riots and Mobs

When a crowd gets completely out of control, as it does in a violent riot or lynch mob, it can do terrible damage. One of the most frightening things about mobs is

that quite normal people can be swept up in them. How can people act in ways that are so out of character? We will see in this chapter that such situations can be very powerful. Somehow the situation leads many people to do things we wouldn't ever think they could do.

### Helping: The Good Samaritan

The book of Luke relates the parable of the Good Samaritan. A man traveling the road from Jerusalem to Jericho is robbed and beaten and left by the side of the road to die. Several people see him but pass him by. Then a Samaritan (a member of a lowly-regarded group) comes upon him, helps, and saves the victim's life.

In 1973 John Darley and Daniel Batson conducted a famous study at the Theological Seminary at Princeton University. The subjects were students at the Seminary, and when each arrived he was told that he would be giving a lecture in another building on campus. Half were told that their talk should be about career alternatives for priests (this was meant to be a neutral topic); the other half were told that their talk should be about the parable of the Good Samaritan. Each group was then divided into three further subgroups that differed only in the instructions they received about how soon the talk was to be.

1. You are already late (high-hurry condition)
2. You should leave now (intermediate-hurry condition)
3. There is no rush (non-hurry condition).

So we have *six groups* in all: two different topics for the talk and three different hurry conditions.

Condition	Low Hurry	Intermediate Hurry	High Hurry
Percent who Helped	63%	45%	10%

Table 23.1: Results of the Good-Samaritan Study

As the subjects made their way to the building where they were to give their talk, each of them passed a man slumped in a doorway. He was coughing, groaning and clearly in need of help. Which groups helped the most?

The topic of the talk didn't make much difference. Furthermore, subjects who scored high on religiosity measures weren't any more likely to help than those who scored low. In fact, the only factor that had much impact on whether a person helped or not was whether he was in a *hurry*: 63% of the subjects in the no hurry condition helped, 45% in the-intermediate hurry condition helped, and only 10% of

those in the high-hurry condition helped (Table 23.1 on the preceding page). The personality or character traits of the subjects surely weren't irrelevant to whether they stopped to give aid. But it was *a feature of the situation*—how rushed the people were—that played the greater role.

### Helping: Kitty Genovese

At 3:20 A.M. on March 13, 1964 Kitty Genovese arrived back at her apartment in Queens after a long night's work. As she walked from her car to her building she was accosted by a stranger who stabbed her repeatedly. Over and over she fell, was stabbed, struggled up, tried to crawl to her doorway, and was stabbed yet again. People in nearby apartments heard her screams, turned on their lights, and watched as the horrifying scene dragged on for almost thirty minutes. But what shocked the nation was that at least 38 people watched the brutal murder and none of them called the police. Not one.

This led the psychologists John Darley and Bibb Latané to wonder about the conditions that would inhibit helping, and in 1969 they conducted an experiment to try to find out. There were three conditions in their experiment. In one condition the subject was alone in the room, in the second the subject was in a group with two other real subjects, and in the third the subject was in a group with two other "subjects" who were actually confederates.

The experiment began normally enough. One or more subjects entered the room and began filling out a questionnaire. But suddenly smoke began coming out of a vent in the room; it certainly looked like something that could be dangerous. When the subjects were alone in the room, 78% of them reported the smoke. When there were three genuine subjects, 38% of the subjects reported it. And when there was one subject together with two confederates who did nothing, only 10% of the subjects reported it.

These results are typical. In 90% of the studies on the matter, a lone bystander is more likely to help than a person in a group. And many studies indicate that your chances of getting help may be best if only one other person around. Why is this—why don't people help when we would expect them to?

### Why Don't People Help?

Why didn't someone call the police as they watched Kitty Genovese's brutal murder? Our first thought might be that they actually enjoyed watching her suffer. But surely all of the people in her neighborhood couldn't have been sadists. In fact, it turns out that when *other people* are present, people in general are more likely to stand by and do nothing. The *situation* inhibits helping.

In the previous chapter we encountered the concept of social proof: people often wait to see what others do in order to determine the appropriate response. If everyone is waiting to see what behavior is appropriate, there may be no response at all. Put yourself in the position of a subject in the smoke study: Maybe the smoke pouring out of the vent is harmless; the other people in this room seem to think so, maybe they know more about such things than I do, and if I go for help I might end up looking like an idiot.

But the Kitty Genovese murder is also a dramatic illustration of something else that is quite common: diffusion of responsibility. If you are the only one present and something needs to be done, you have to do it if it is to be done at all. But if there are several people around, maybe someone else will take action, and you'll be off the hook. **Diffusion of responsibility** occurs when a number of people are present, and the responsibility diffuses or radiates throughout the group, so that no one feels particularly accountable. In situations like this, people are less likely to help.

### Conformity Revisited

#### The Autokinetic Effect

In the previous chapter we learned about Sherif's conformity studies involving the autokinetic effect. The autokinetic effect is a perceptual illusion in which a stationary point of light seems to move and then disappear. Different people think it moves different distances (from an inch up to several feet). But when Sherif's subjects worked in groups, they converged on group norms that were very strong and persistent. Although the group members were not aware of it, there were powerful pressures in the situation that led them to conform.

#### Ash's Lines

In Ash's conformity studies subjects went along with the obviously incorrect answers of the confederates far more often than anyone would have predicted. Indeed, Ash himself thought Sherif had exaggerated the tendency for people to conform, and he originally began his study to show the limits to conformity. But Ash had created a very powerful situation—one full of *conformity pressures*—that exerted a strong influence on his subjects' behavior. And this occurred even though the group members didn't know each other and wouldn't be interacting in the future.



### Obedience Revisited

In one of Milgram's experiments about 65% of the subjects went all the way, administering shocks of 450 volts, and *all* of them went to at least 300 volts before breaking off. Moreover, the rate was about the same for men and for women.

*No one*—students, laymen, psychologists, psychiatrists—expected so much obedience. Nobody's predictions were anywhere close. Many experts said that only one person in 1000 would go all the way to 450 volts. Many were confident that people would defy the experimenter and quit at about 150 volts, the point where the learner first asks to be released. How could all these predictions have been so wrong?

The easiest explanation is that they were cold-hearted people who didn't mind inflicting pain. Indeed, this is exactly how subjects in a later study (by Arthur Miller and his coworkers) viewed Milgram's subjects. But there were hundreds of subjects; what are the chances that Milgram just happened to find a group of several hundred people, two thirds of whom were sadists? It's not likely.

In any case, the subjects did not enjoy administering the shocks; most of them were very uncomfortable about what they were doing (even when they kept right on going), and many remained shaken by the experience afterward. The experts had based their predictions on the internal (dispositional) characteristics of people; in their view, only people who were abnormal or deviant in some way would go on and on with the shocks. They didn't realize that Milgram had created a very powerful *situation* in which obedience was very easy and very likely to occur.

This is not to deny that people are different. Some students did stop to help the person in the doorway; some subjects did defy the experimenter and quit delivering shocks. The *point* is simply that *many more* people were carried along by the power of the situation than anyone would have expected.

### Prisoners and Guards

On Sunday morning, August, 17, 1971, nine young men were picked up without warning at their homes by the Palo Alto police. They had been drawn from a group of about seventy men who had answered an advertisement in the local paper offering \$15 to participants in a two-week study on prisons. After interviews and psychological screening, the group was narrowed to about twenty five, and these people were randomly assigned to play the role of prisoner or guard.

The nine men arrested that Sunday morning were those who had been randomly assigned the role of prisoner. They were driven to the local Police Station, booked, fingerprinted, blindfolded, and taken to a simulated prison in the basement of the psychology building at Stanford University. Meanwhile, those assigned the role

of guards were given uniforms and instructed that their task was to maintain order (without using violence).



The subjects were part of a study on roles and behavior conducted by the Stanford psychologist Philip Zimbardo and his coworkers. After an initial rebellion by the prisoners, the guards quickly gained control, and soon got completely into their role as guards. They taunted, humiliated, and degraded the prisoners, making them do pushups or clean out toilet bowls with their bare hands when they didn't obey. They began treating the prisoners like they weren't real human beings.

The prisoners also got into their role as prisoners, becoming listless, subservient, and suffering from stress (some had to be released early because they were cracking under the pressure). In fact their reaction was so severe that the experiment had to be called off before the end of the first week.

The subjects were assigned *randomly* to play the *role* of prisoner or guard. But the situation felt so real, with uniforms, bars on the

cells, and the other props of a real prison, that the subjects quickly adopted their roles all too well. In a very short time, normal people were transformed into sadistic guards or passive victims. Zimbardo and his coworkers had created a very powerful situation in which people fell into predetermined roles despite themselves. If six days in a setting that everyone knew was “just an experiment” had this effect, what effects might an even more powerful situation (like a real prison) have?

### **Brown Eyes vs. Blue Eyes**

Even young children are not immune. Jane Elliott was a third-grade teacher in the small Iowa town of Riceville. Her students had little exposure to minority groups, so she decided to let them learn first hand. One day, in the late 1960s, when her students arrived for class Elliott informed them that brown-eyed children were smarter and better than blue-eyed children and so they should be treated better. The superior, brown-eyed students were then given various privileges while the inferior blue-eyed students were subjected to demeaning rules that underscored their inferior, lowly status.

Well before the end of the day the brown-eyed students were discriminating against their blue-eyed former friends: they fought with them, ostracized them, and suspected them of underhanded behavior. Meanwhile the blue-eyed students became angry, demoralized, and withdrawn.

The next day Elliott told the students that she had made a mistake; it was actually the blue-eyed children who were superior. The situation then replayed itself with the blue-eyed children engaging in ready and often hostile discrimination. The third day the class discussed the implications of what they had been through. In 1992 before a huge television audience on the Oprah Winfrey show Elliott carried out the experiment, with similar results, using adults as subjects.

## 23.2 The Fundamental Attribution Error

### 23.2.1 Explaining why People do what they Do

We are often more interested in other people and what makes them tick than in anything else. *Why* do they do the things that they do? What led several hundred people at Jonestown to happily commit suicide? Why did so many of the subjects in Stanley Milgram's famous experiments on obedience to administer (what they thought were) severe shocks to the learner? Why did so many people do nothing while Kitty Genovese was murdered? Such questions also arise closer to home. Why did Sally give Wilbur that weird look when he said that they should go out again soon? In fact, we often have occasion to wonder why *we* do some of the things that we do: why in the world did I say *that*?


#### Patty Hearst

On Friday, February 4, 1974 Patty Hearst, the eighteen year old daughter of a wealthy San Francisco publishing family, was kidnapped by a terrorist group calling themselves the Symbionese Liberation Army (the SLA). She was abused and tortured and kept—bound and blindfolded—in a closet for 57 days. It is not surprising that she was terrified.

What is surprising is what happened next. Hearst began to identify with her captors. She renamed herself 'Tania', and carried a machine gun into a San Francisco bank and held it on the customers while other members of the SLA robbed the place. Even twenty months after her rescue, she continued to defend the views of the group. Why did she do this? There was nothing in her past to indicate that she would have any sympathy with a radical group like the SLA, and even the people who knew her best couldn't understand it.

There probably isn't any simple explanation for her actions, but there are two

general *types* of answers. First, perhaps she was one of those rare people who gets swept up in such things; she had a weak character and wasn't strong enough to resist.



This may well be part of the story, but a very different *sort* of answer is possible. It may be that many kidnapping victims begin to identify with their kidnapers after a certain period of time. In fact there is a name for this phenomenon: it is called the 'Stockholm Syndrome' (after a 1973 incident in Stockholm in which robbers held four people captive in a bank vault for six days; after several days the victims began to establish a bond with their captors). Some psychologists claim that such behavior is a not-uncommon attempt to cope with the uncertainty and terror in the situation. But we won't be concerned here with which explanation of Hearst's behavior is correct (quite possibly both get at part of the truth). Our interest is in the two, quite different *types* of explanations illustrated in this paragraph.

### Internal vs. External Causes

We can try to explain Hearst's behavior by citing "internal" causes (her character traits, e.g., being weak and impressionable) or by citing "external" causes (the fact that many people in a terrifying situation like hers start to identify with their captors in order to cope with their terror). More generally, we can divide the causes of peoples' actions into two sorts:

**Internal Causes:** Causes "inside" the person: his or her personality traits or dispositions, attitudes, values, desires.

1. John returned the lost billfold because he is *honest* and *helpful*.
2. Wilma yelled at John because she was *angry*
3. Wilbur gave me that weird look because he's an extremely *creepy* individual.

**External Causes:** Causes "outside" the person: features of the *situation* in which the person acts.

1. John said that line 2 matched line A because of the strong *social pressure* exerted by the other people in the experiment.
2. Wilma didn't help the victim because the *situation* was one in which it was unclear whether the victim needed help and no one else at the scene seemed to think that he did.
3. Wilbur followed the leader's orders, but *anyone else* would have done the same thing in those *circumstances*.

All actions take place in some context or situation, and *both* a person's internal states (dispositions, attitudes, etc.) *and* features of the situation (e.g., the presence of others, the commands of an authority figure) play a role in determining what he does in that situation. It is *never* the case that internal causes aren't important. But as we will see, people strongly overestimate the strength of internal causes while underestimating the strength of external ones.

### The Fundamental Attribution Error

All these cases—mobs, helping, conformity, obedience, the prison study—illustrate *the power of the situation*. But we tend to underestimate this power. This is such a large and common bias in our reasoning about other people that it has been given a name: the fundamental attribution error. **The fundamental attribution error** occurs because of our strong tendency to overestimate the significance of internal causes and to underestimate the power of external (situational) causes.

**Fundamental attribution error:** underestimating the power of the situation to influence behavior

The fundamental attribution error gets its name from the fact that we often make this mistake when we are trying to *attribute* a person's actions to causes of one kind or another. For example, we commit this error if we focus too much on whether a would-be helper is a helpful sort of person (thus attributing their behavior to an internal cause) while overlooking features of the situation like the fact that other people are present (thus ignoring external causes).

### The Fundamental Attribution Error in the Laboratory

People commit the fundamental attribution error in the real world and in the psychological laboratory. In various studies subjects listen to someone give a speech in which they read an essay that was actually written by the experimenter. The speech defends some cause like the legalization of marijuana.

Even when subjects are *told* that the person was required to give the speech and that it may not reflect her true view, they are strongly inclined to believe that it really does reflect her views. In this case the subjects do not take the situation adequately into account (the other person was *required* to give this speech, and almost anyone would do the same thing in such circumstances).

In another experiment, Lee Ross and his collaborators had subjects play a quiz game. It was clear to everyone involved that subjects were *randomly* assigned to be either the questioner or the contestant. The questioner was instructed to devise ten difficult factual questions based on her own knowledge which she was then to ask the contestant. The questioners were at a very great advantage since *they* picked the questions based on their own background and expertise (which the contestants were unlikely to share).

But despite this clear situational advantage, observers, questioners, and even the contestant themselves rated the questioners as more knowledgeable and intelligent than the contestants. People underestimated the power of this situation—the advantage of those who got to make up the questions—and overestimated the extent to which the contestants' behavior reflected their traits or characteristics (like being intelligent or knowledgeable).

Another way to see the point is that the results would have been quite different if the questioner and the contestant had exchanged roles; then the people who actually seemed smarter would have seemed less intelligent, and vice versa. The situation is such that the questioner—whoever it is—will look better, but people tend to overlook this fact. When they do, they commit the fundamental attribution error.

### **The Fundamental Attribution Error in the Real World**

When we first hear about the thirty eight bystanders who watched Kitty Genovese murdered or the people in the Milgram study who administered ever-greater shocks, we are first inclined to think they are uncaring, cruel or sadistic. When we do so, we attribute their behavior to *internal causes* (their uncaringness, cruelty or sadism). We overlook the fact that the situations are very powerful and that many people—perhaps even us—would act the same way in those situations. When we do this, we commit the fundamental attribution error.

The fundamental attribution error is also encouraged by the belief that people have relatively stable traits that strongly influence how they will behave in a wide range of settings: Wilbur is honest, and he would behave in an honest way in almost any circumstances. But it turns out that people's traits aren't as robust as we usually assume. There is not as much consistency in people's behavior from one type of situation to another as we commonly suppose.

### **What The Fundamental Attribution Error does *Not* Mean**

Before proceeding it is important to note two things that do *not* follow from the fundamental attribution error. First, the claim is not that every one is the same.

People *do* differ, and these differences help account for why they do the things that they do. If virtually everyone in a situation would do the same thing, e.g., eat grasshoppers when an experimenter pressures them to, then the fact that a particular person ate some grasshoppers doesn't tell us much about her.

On the other hand, if somebody does something that most people would *not* do in that same situation, their action *does* tell us something about them. For example, most people in Clinton's position would not have gotten involved with Monica Lewinsky, so the fact that he did tells us something *about him*. The point is not that we should never attribute behavior to internal causes but that we tend to *overattribute* it to internal causes.

Second, the fact that situations are more powerful than we often suppose does not mean that people are not responsible for what we do ("It's not my fault: the situation made me do it"). Some people help even when others are present. Some people refuse to go on shocking an innocent victim. Some Europeans hid Jews during World War II in the face of strong social pressures and grave physical dangers.

Indeed, the hope is that by learning about the power of the situation we will be better at *resisting* that power. Learning about the frequent failure of people in groups to help someone in need should make it easier for us realize the importance of stopping to help. And learning about the Milgram experiment should make it easier for us to stop and ask, when someone in authority tells us to do something that seems questionable, whether we should comply.

### Consequences of the Fundamental Attribution Error

The fundamental attribution error is a very common bias in our reasoning about other people, and it can lead us astray in several ways.

1. It leads us to think that they are *more consistent* than they actually are.
2. It leads us to think that we can do a better job of *predicting* their behavior on the basis of their traits than we can
  - We would often do better basing our prediction on our knowledge of the situation.
3. It leads us to think that we have a *better understanding* of human behavior than we do.

But the fundamental attribution error also suggests some more positive lessons. It is important to raise people with good character. But since behavior is more strongly influenced by situations than we often suppose, it is *also* important to design social settings and situations in a way that is likely to bring out the best in people, rather than the worst.



### 23.3 Actor-Observer Differences

**Actor-observer asymmetry:** we tend to see other people's behavior as internally caused, but to see our own as externally caused

We often explain other people's behavior by citing internal causes like their beliefs and attitudes and traits. John helped the old lady carry her groceries because he's a caring and helpful person. But how often do we explain our own actions this way? How natural would you find it to say: "I helped the old lady carry her groceries because I'm caring and helpful person."

Of course we might not *say* this because it sounds so immodest. But how often do we even *think* of our own actions in this way? We are much more likely to say (and think) that we helped the old lady because she looked frail and in need of help. In doing so, we cite features of the situation (a frail old lady needing help) rather than internal causes (I'm such a helpful person).

This asymmetry in how we think about the actions of others and our own actions is known as the **actor-observer difference** (or the **self-other difference**). We tend to see *other peoples'* actions as having internal causes (John is helpful) but we see our *own* actions as having external (situational) causes (she needed help). This phenomenon gets its name because the agent or *actor* (in this case John) sees his actions are largely influenced by the situation. But when John *observes* others, he sees *their* actions as largely influenced by their traits and other internal states.

The actor-observer difference amounts to a bias in our reasoning about people's actions, both our own and those of others we observe. But it does mean that we are less susceptible to the fundamental attribution error when we try to explain our own actions than when we try to explain the actions of others.

### 23.4 Special Cases

There are certain special cases where we tend to give situational explanations sorts of behavior, and other cases where we tend to give dispositional explanations. We conclude with three important examples.

#### Blaming the Victim

Earlier we discussed the just-world hypothesis. We tend to see people as getting pretty much what they deserve, and when someone suffers in a way that is not the result of obviously bad luck, we tend to think they must have brought it on themselves. They suffer their misfortune because of the *sorts* of people they are. When we reason in this way we are giving a *situational explanation* of their behavior.



### Ultimate Attribution Error

Many people tend to give dispositional explanations of the failures or the negative behavior of members of groups they don't like: "It's no wonder Wilbur did poorly on the exam; he's an Okie, and Okies are dumb." By contrast, we tend to explain their successes and positive behavior in situational terms: "He was just lucky," "She must have gotten some special break". We will return to this matter in more detail in Chapter 26.

### Self-Serving Biases

We are more likely to attribute *our own* good or successful actions to internal causes and our bad or unsuccessful ones to external causes. There are actually two biases here. The **self-enhancing bias** is the tendency to attribute successful outcomes to our own abilities, and the **self-protective bias** is the tendency to attribute unsuccessful outcomes to the situation. There appears to be a stronger tendency to take credit for our good actions than to blame our failures on the situation, although the issue is a difficult one to study because people may not report their true feelings when discussing themselves and their own actions.

Biases in how we think about ourselves are related to the *Lake Wobegon effect* (p. 164). There we learned that a large majority of people think that they are above average in a variety of ways, and only a very small percentage think that they are below average.

For example a survey of a million high-school seniors found that 70% rated themselves above average in leadership skills while only 2% felt they were below average. And *all* of them thought that they were above average in their ability to get along with others. Most people also think of themselves as above average in intelligence, fairness, job performance, and so on. They also think they have a better than average chance of having a good job or a marriage that doesn't end in divorce.<sup>1</sup>

### Conclusion

A great deal of our thinking in daily life involves thinking about people, others and ourselves, trying to understand, explain, and predict their actions. Two of the more

---

<sup>1</sup>For an excellent discussion of the power of the situation and related matters like the fundamental attribution error see Lee Ross and Richard Nisbett's *The Person and the Situation*, McGraw Hill, 1991. This book also contains references to many of the studies described above. Jane Elliot describes her study with her third graders in "The Power and Pathology of Prejudice," in P. G. Zimbardo & F. L. Ruch, editors, *Psychology and Life*, Glenview, IL: Scott, Foresman, Diamond Printing, 9th/ed. Further references to be supplied.

robust findings in recent psychology concern such social cognition or reasoning. First, we tend to commit what has come to be called the fundamental attribution error and, second, actors and observers view the causes of behavior differently. In this chapter we have seen the various ways in which we are susceptible to these biases and a number of the ways they lead to suboptimal reasoning.

### 23.5 Chapter Exercises

1. Several studies suggest that people in Asian societies like Taiwan are less likely to commit the fundamental attribution error than people in the U.S. or Western Europe. What does this claim mean? What might explain this?
2. Even when people are expressly told that an essay's author was forced to defend a particular position, they tend to attribute that position to the author. Why?
3. Wilbur arrives thirty minutes late to pick up Wilma for their first date. He tells her that he was detained on the phone with his mother, and then he had to stop and get money, then he had a flat.
  1. How does Wilbur probably view this situation?
  2. How does Wilma probably view this situation?
  3. How could Wilma's view, in conjunction with the primacy effect, affect her later views about Wilbur?
4. Describe a real-life situation (one not mentioned above) that involves actor-observer differences. Explain what is going on with our reasoning when such differences occur.
5. What implications do the themes of this chapter have for self-knowledge, for the degree to which we can understand who we are and why we do the things we do?

# Chapter 24

## Reasoning in Groups

**Overview:** Being able to work in groups is important in today's world. Many projects are carried out by teams, numerous decisions are made by committees, and most people's jobs require them to work as a part of a group. Juries, parole boards, city councils, corporate board meetings all reason and make decisions. And families are groups, ones whose decisions affect us quite directly. In this module we will consider several features of group reasoning and decision making. There is great variability among groups, so we can't expect any simple, blanket conclusions. But we will see that although groups have their virtues, they are susceptible to several sorts of biases and we need to guard against them.

### Contents

---

<b>24.1 Group Reasoning</b> . . . . .	<b>486</b>
<b>24.2 Social Loafing</b> . . . . .	<b>486</b>
<b>24.3 Group Dynamics and Setting the Agenda</b> . . . . .	<b>487</b>
24.3.1 Heuristics and Biases in Groups . . . . .	487
24.3.2 Out-group Homogeneity Bias . . . . .	487
<b>24.4 Group Polarization</b> . . . . .	<b>488</b>
<b>24.5 Group Accuracy</b> . . . . .	<b>489</b>
<b>24.6 Groupthink</b> . . . . .	<b>489</b>
<b>24.7 Successful Groups</b> . . . . .	<b>489</b>
24.7.1 Groups in the Classroom . . . . .	489
<b>24.8 Safeguards</b> . . . . .	<b>491</b>

---

24.9 Chapter Exercises . . . . . 491

---

## 24.1 Group Reasoning

Many policies are fashioned or overseen by legislative bodies, advisory panels, committees, coalitions, boards, or other groups. Under some conditions the cognitive shortcomings of group members can be attenuated by the group, but under some conditions they are accentuated. Groups are also susceptible to biases of their own, including polarization (making a more extreme decision than group members would individually), out-group homogeneity (seeing other groups as more homogeneous than they really are), and such nebulous but real afflictions as Janis' "group think" (which occurs when a group accepts feels that it must be right and neglects to make a reality check).

We often *feel* that groups are more likely to arrive at more balanced, reasonable conclusions and decisions than individuals working alone. Other things being equal, we tend to suppose, groups have the following advantages:

1. They typically have more information than single individuals.
2. More viewpoints are likely to be represented.
3. Problems that one person might overlook are more likely to be noticed.
4. They are likely to take fewer risks and make less extreme recommendations.

And indeed groups often do a better job than individuals working alone. But not all of the above points are true of all groups, and groups also exhibit various biases and weaknesses.

## 24.2 Social Loafing

**Social loafing:** members of a group often do less than when working alone

Members of a group often do less work, and do it less well, than they would do if they were working alone. This phenomenon is called **social loafing**. One reason for social loafing seems to be a *diffusion of responsibility*. Each member of a group feels less responsibility and accountability for the work than they would if they had the sole responsibility. This is confirmed by the fact that social loafing can be reduced if each member of a group has a specific task or if each member is accountable for the work that they do.

Many college classes nowadays feature group projects. Most of the students in this class are not enthralled with group work, and the most common objection is

precisely that they promote social loafing, so that some people do more than their share of the work. One way to reduce social loafing is to assign each member of the group a specific role and to grade them on their contribution to the group. Social loafing is relevant in understanding the dynamics of groups and in explaining the behavior of their members. Both of these are related to reasoning, but in this chapter we will focus on several topics that involve reasoning even more directly.

## 24.3 Group Dynamics and Setting the Agenda

Many group deliberations exhibit a similar structure. A range of options are proposed and discussed. At some point an option arises that no one strongly objects to (even if they don't like it very much). At this point further options are not well received, and some version of this proposal has a good chance of being accepted. In other words a group tends to focus on ideas that happen to be brought up early in the discussion and to give most of their weight to preferences that are expressed relatively early.

This means that the order in which options are introduced can affect a group's decision (later we will see that the order in which things are voted on can determine which one will win). This is so, because groups often have a bias to minimally acceptable solutions that come up relatively early in discussion.

### 24.3.1 Heuristics and Biases in Groups

Throughout this course we have seen that individuals are susceptible to various fallacies and biases. It would be nice if groups were less affected with such maladies, but research suggests that this is not always the case. Groups can commit the conjunction fallacy, rely too heavily on inferential heuristics, and ignore information about base rates. Moreover, just as individuals are often guilty of self-serving biases, groups are often guilty of group-serving biases.

### 24.3.2 Out-group Homogeneity Bias

But new sorts of biases enter the picture when we turn to groups. Perhaps the most important is the tendency to see other groups (especially "out-groups"; e.g., people of other races, religions, countries, sexual-orientation) as *more homogeneous* than they really are. In other words, groups tend to see themselves as quite varied, whereas members of other groups are thought to be much more similar to one another. Not only are the members of out-groups seen as more similar than they really are. They are often thought to share stereotypical attributes.

**Out-group homogeneity bias:** tendency to see out-groups as more homogeneous than they really are

There are probably various reasons for this. Often in-group members have limited interaction with those in the out-group, and so they don't see the wide variations among members of that group. In this case, people rely on a small and biased sample of the out-group based on limited interactions, news reports, or conventional "wisdom" about members of the out-group. As we will see in a later chapter, this bias about other groups can promote stereotypes, which in turn lead to prejudices and ethnocentrism. If the members of a group seem reasonably similar, then we will be more inclined to generalize from the behavior of a small sample of them than if we thought of the group as more heterogeneous.

## 24.4 Group Polarization

Group decisions are often more extreme than the decisions their members would make if they were acting alone. The first example of this that was studied by social scientists was the risky shift.

**Group polarization:** group decisions are often more extreme than the decisions their members would make acting alone

The **risky shift** occurs when people who take part in a group discussion are willing to support riskier decisions than they would individually, before the group discussion. This has been found in numerous studies. Many of them employ a questionnaire in which a person or group is asked what the probability for success would have to be in order to recommend the more risky (but more profitable) of two options. For example, how likely does the prospect for success need to be to warrant switching to a job with higher chances of advancement but also higher chances of failure?

Many decisions in the real world are made by groups, so the question naturally arises whether these are more risky than decisions made by individuals would have been. (Can you think of any examples?)

More recent research has also found evidence of a **conservative shift**. This is just the opposite of a risky shift; it occurs when people in a group decide to support less risky decisions than they would individually. Both shifts are examples of a more general process of **group polarization**, which can lead either to riskier or to more cautious decisions, depending on the initial views of group members. A risky shift tends to occur when the views of group members already incline towards being risky. A conservative shift tends to occur when they already incline towards being conservative.

Group polarization can affect attitudes as well as decisions. For example, it has been found that students with relatively little racial prejudice became less prejudiced after discussing racial issues, whereas students who began with prejudices emerged from the group discussion with even stronger prejudices. Similar results have been found on the issue of women's rights and many other topics.

## 24.5 Group Accuracy

How accurate are the conclusions groups reach? It depends on the problem and the makeup of the group. But in many cases when a task (like a mathematical problem) has a clear right answer, groups tend to do better than the members do on average, though not as well as the best members of the group.

## 24.6 Groupthink

**Groupthink** is an impairment in decision making and judgment that can occur in highly cohesive groups with a strong, dynamic leader. Group members isolate themselves from outside information, try to please the group leader, and agree on a decision even if it is irrational. They tend to have feelings of invulnerability, and a strong sense that they are right. Such a group is in strong need of a “reality check.” But since there is so much pressure inside the group to agree, it is not likely to come from inside the group. And since the group often feels isolated, sometimes even persecuted, it is not likely to listen to voices from outside the group.

**Groupthink:** tendency for groups to reach uncritical decisions because of group pressures to agree

History is full of such cases. Recent examples include David Koresh and the Branch Davidians in Waco, Texas, Jim Jones and his followers in Jonestown, and the members of Heaven’s Gate.

More importantly, groupthink can occur on less extreme scales. Often top governmental decision makers are part of a small, closed team, and they too can fall prey to groupthink. For example, the America’s attempts to overthrow Castro in Cuba led to the Bay of Pigs fiasco in 1961. In view of the lack of preparation, this was quite predictable, but none of the top officials in Kennedy’s government saw the disaster coming. Even after their landslide victory over George McGovern in 1972, the Nixon Whitehouse felt under siege, going so far as to construct an official “Enemies List.” Their schemes led to Watergate and Nixon’s eventual resignation.

## 24.7 Successful Groups

Although groups, like people, can be biased, there are several things that promote good group reasoning. The points here apply to many groups, but since you are probably doing group work in some of your classes, we will focus on that.

### 24.7.1 Groups in the Classroom

More and more classes include group projects. There is little evidence that this improves learning, although it probably does help students reason in a social context.

Some students dislike group work because they feel they end up doing more than their share while others, the social loafers, get a good grade despite doing little work. But even if some students do more of the work, they do gain something from it. In the chapter on memory we distinguished between recognition and recall. With the former we can recognize something when we see it again, but we couldn't have said much about it if we hadn't. If we can recall something, we have a better grip on it. And, continuing down this path, if we can explain something to others we have still a better grip. Being in a situation where we have to explain things promotes the depth of our own learning and understanding.

If a group project involves a newly-formed group that will only be together for one class meeting, the prospects for useful discussion and learning are not great. The group is not likely to be cohesive, and in this setting people tend to avoid disagreement. Hence, even if someone proposes an idea or solution that many of the group members think is flawed, they often are unlikely to say so. Moreover, others may be reluctant to advance their own ideas, because they are worried about the reception they will get. This lack of open discussion makes it much harder to critically evaluate the ideas that are proposed, and it may prevent some good ideas from being proposed at all.

If a group is to make head-way with a difficult or controversial issue, an open expression of alternative points of view followed by a frank discussion of them is needed. If the group has met several times before and its members have gotten to know one another, there is likely to be a higher level of *trust*, and this will make it easier for each member to express her own ideas and to candidly evaluate the ideas of others.

Things also work much better when the group have very specific goals. In the absence of highly specific tasks, discussion often becomes somewhat aimless or drifts off into small-talk.

Groups also work better when they are given quick feedback. This isn't surprising, because feedback is always important in learning; if you don't know how you are doing, it's difficult to make progress. But it may not be so obvious that groups do better when this feedback includes a comparison of the their work with the work of other groups. Such comparisons foster group cohesiveness and effort, because group members then see themselves in competition with other groups. They are then less likely to compete against one another and more likely to work together to compete against other groups.

Finally, groups do better if there are rewards for success. We all are most likely to put our time and energy into projects that have a definite payoff.

All of these points demonstrate once again the power of the situation. How well people do in groups depends on



1. Whether they have a cohesive group in which members trust each other.
2. Whether they are given a definite task with a clearly defined goal.
3. Whether they get quick feedback, including feedback on how their work compares with that of other groups.
4. Whether they are rewarded for good work.

Although we have been discussing groups in the classroom, most of these points are quite general. Some groups (e.g., juries) meet only once; others (e.g., a city council) meet repeatedly over time. Some (e.g., juries) have a clearly defined task; others (e.g., a task force to solve some social problem) have a more nebulous task.

## 24.8 Safeguards

Once again an atmosphere that fosters dissent is important, and the group leader should encourage people to air objections. Many group biases are less likely if group members are allowed to express disagreement, since the group will be exposed to alternative points of view. It will be easier to break the grip of a particular way of looking at the relevant issues. It can even be useful to have someone play the role of devil's advocate. Sometimes it is also helpful to first discuss the issues in subgroups (to encourage development of alternate points of view). And seeking the views of people from outside the group also provides another perspective.<sup>1</sup>

## 24.9 Chapter Exercises

1. You are teaching a class and want to assign group projects. Answer each of the following questions *in detail*, and give examples.
  - 1a. What sort of feedback would you give the groups?
  - 1b. What steps would you take to make sure they knew how well they were doing in comparison to the other groups in the class?

---

<sup>1</sup>The account of the common dynamics of group decision making is described in R. Hastie, S. D. Penrod, and N. Pennington, *Inside the Jury*, Cambridge, MA: Harvard University Press, 1984. Also relevant is L. R. Hoffman and N. Mair, "Valence in the Adoption of Solutions by Problem-Solving Groups: Concept, Method, and Results," in L. R. Hoffman, ed., *The Group Problem Solving Process: Studies of a Valence Model*, NY: Praeger, 1979. For the effects of group polarization on racial prejudices see D. Myers and G. Bishop, "Discussion Effects on Racial Attitudes," *Science*, 169, 1970; 778-779. For a good discussion of groups in the classroom see L. Michaelsen, R. Black, and L. Fink, "What Every Faculty Developer Needs to Know about Learning Groups," in L. Richlin, ed., *To Improve the Academy: Resources for Faculty, Instructional and Organizations Development*. Stillwater, Ok: New Forums Press Co., 1996.

1c. What steps would you take to make sure groups were rewarded for their efforts?

1d. How could you give rewards in a way that didn't unduly penalize those who did more work or unduly reward those who did less?

1e. More importantly, how would you encourage everyone to do their fair share in the first place?

2. Suppose that you learn that some important policy decision was made by a large group. How might that affect the way you think about it?

3. What do you think explains the risky shift? What role, for example, might diffusion of responsibility play here?

4. Give an example from your own experience of the risky shift (also called the choice shift) and explain the things that you think brought it about. What problems might a risky shift cause and how might we try to minimize its dangers?

5. Give an example from your own experience of group polarization and explain the things that you think brought it about. What problems might a group polarization cause and how might we try to minimize its dangers?

6. Give an example, either from your own experience or from your reading, of "groupthink," and explain the things that you think brought it about. How might we try to minimize the tendency for people to engage in groupthink?

7. Note some concrete ways in which we might try to reduce the various dangers inherent in group thinking. How practical are they?

# Chapter 25

## Stereotypes and Prejudices

**Overview:** The harmful consequences of stereotypes, prejudice, and discrimination are all around us. In this chapter we focus on the ways fallacies, biases, and other flawed patterns of thinking help foster and maintain prejudices and stereotypes. It would be too much to hope that clearer thinking alone would eliminate such problems, but it would be a step in the right direction.

### Contents

---

<b>25.1 Consequences of Prejudices and Stereotypes</b> . . . . .	<b>494</b>
25.1.1 Stereotypes, Prejudice, and Critical Reasoning . . . . .	495
<b>25.2 Prejudice</b> . . . . .	<b>495</b>
<b>25.3 Stereotypes</b> . . . . .	<b>495</b>
25.3.1 Schemas . . . . .	495
25.3.2 Stereotypes as Schemas . . . . .	496
<b>25.4 Discrimination</b> . . . . .	<b>497</b>
25.4.1 Overt vs. Modern Racism . . . . .	498
25.4.2 In-groups vs. Out-groups . . . . .	498
<b>25.5 Features of Stereotypes</b> . . . . .	<b>500</b>
25.5.1 Homogeneity . . . . .	500
25.5.2 Resistance to Change . . . . .	501
<b>25.6 Cognitive Mechanisms</b> . . . . .	<b>502</b>
25.6.1 Levels of Generality . . . . .	503
<b>25.7 Flawed Reasoning and Prejudice</b> . . . . .	<b>504</b>
<b>25.8 Responses</b> . . . . .	<b>515</b>

25.8.1 Remedies and Reasoning . . . . .	515
25.9 Chapter Exercises . . . . .	517

## 25.1 Consequences of Prejudices and Stereotypes

The harmful effects of stereotypes, prejudice, and discrimination are all around us. At their worst they can lead to institutionalized discrimination, slavery, and even genocide (like the Nazi's killing of over six million Jews or the recent "ethnic cleansing" in Kosovo). But milder forms of prejudice have an adverse impact every day on people all around us. They can affect every aspect of a person's life, from her opportunities as a child to her dignity and self-esteem. They also are bad for those who hold them; going through life hating people simply because they are members of a particular group is not a rewarding way to live.



The catalogue of atrocities resulting from prejudice and discrimination is virtually endless. According to our best scientific estimates several centuries ago there were five to ten million Native Americans living in what is now the Continental United States and Canada; by 1900 the number had fallen below 400,000. Over three thousand African Americans were lynched in the United States from the 1880s to the 1960s. And the list goes on and on.

Although violent hate crimes still occur and there are still organized groups like the Klu Klux Klan and Neo-Nazi organizations whose chief purpose is to promote prejudice, blatant expressions of prejudice are now much less common, but more subtle versions of discrimination still blight the lives of many.

A finding recently reported in the *New England Journal of Medicine* (Feb. 25, 1999) illustrates the point. In a study involving actors (professing to list their symptoms) and 700 real doctors, blacks, particularly black women, were less likely than whites to receive proper testing for serious heart disease—even when doctors believed that the black women had just as high a likelihood of benefiting from the test.

Few (if any) of the doctors were blatant racists. But the results of their decisions on this life and death matter would still have favored whites.

### 25.1.1 Stereotypes, Prejudice, and Critical Reasoning

Stereotypes and prejudices have many causes and meet many needs. Some of the causes involve emotions and feelings that aren't easily affected by evidence or reasoning, and so attempts to change things will certainly require more than careful and critical thinking.

But faulty reasoning *does* help foster and maintain stereotypes and prejudices. Since our concern in this book is with critical reasoning, we will focus on how the fallacies, biases, uncritical use of various heuristics and other flawed patterns of thinking studied in earlier chapter help support stereotypes and prejudices. In short, this chapter is an application of our earlier work to a very important—but very difficult—social topic.

Although it would be too much to hope that clearer thinking alone would eliminate prejudice and discrimination, it would be a step in the right direction. Among other things, it would make it more difficult to defend stereotypes with bad reasoning; it would also help us think about these topics in a way that isn't so colored by emotions that we don't think about them carefully.

## 25.2 Prejudice

A **prejudice** is an attitude toward a group and its members. The members are targeted simply because they belong to the group. Prejudices typically involve beliefs (stereotypes) about the group *and* emotions or feelings (ranging from mild dislike to extreme hatred). Prejudices are often held for emotional reasons or to help rationalize actions like exploitation of a minority group. But people with prejudices often defend their views—which of course they don't think of as prejudices—with reasons and arguments. And these *can* be rationally evaluated.

## 25.3 Stereotypes

### 25.3.1 Schemas

As we saw earlier (Ch. 7.5.2), there is now considerable evidence that people have well-organized packets of generic knowledge about many things, including librarians, picnics, graduate student offices, visits to restaurants, first dates, and so on. These packages of information are called *schemas*.

For example, the picnic schema includes having a meal outdoors, eating off paper plates, taking a bottle of mustard, sitting at a picnic table, and so on. Not all of these things are required for a picnic, but most of them will be true of your typical picnic. So once we learn that Wilma and Wilbur went on a picnic we can reasonably infer that most of these things probably occurred; e.g., they probably took a bottle of mustard and ate off paper plates.

Schemas are useful because they help us organize our knowledge, they allow us to automatically fill in many details, and they structure our expectations. A little information may activate it, and we can then use the generic knowledge in the schema to quickly draw further *inferences* about the situation or thing.

For example, when we hear that Wilbur packed a bottle of mustard in a basket, this is likely to activate our picnic schema. And we can then use it to draw inferences about what Wilbur is up to; we would expect him to be eating outdoors, for example. Similarly, your schema for a graduate student office probably includes having books in it, so it is natural to expect that it will or to remember that it did.

We couldn't get along without schemas. They help us categorize things and people, which makes it easier to explain and predict what they do. Even schemas that are a bit inaccurate are often useful. But some schemas are *so* inaccurate that they often lead to bad consequences.

### 25.3.2 Stereotypes as Schemas

A **stereotype** is a schema about the members of a certain group. It includes the character traits, physical traits, and other features commonly associated with the group. When we think of a picnic we are likely to think that people ate off paper plates, or to remember that they did (even if they didn't), because this is part of our schema for a picnic.

Similarly, research shows that many prejudiced whites have a stereotype of Afro-Americans males as athletic, hostile, and lazy. So when someone who has a schema like this learns that Wilbur is an Afro-American, they are likely to infer (virtually automatically) that he is probably hostile. A stereotype needn't involve negative feelings about the stereotyped group, but when it does, it leads to prejudice.

Stereotypes are dangerous because they are (i) typically inaccurate, (ii) attribute the same features to virtually all members of a group (even though there are usually large differences among group members), and (iii) are often very resistant to change. We will examine these features below, but first it will be useful to consider some examples of stereotypes.

### Examples of Stereotypes

It is often uncomfortable to discuss stereotypes and prejudices in class. For one thing, some of the people in the class are likely to be members of targeted groups. But open and respectful discussion is necessary if we are to think about these issues in a clear way. If we want to see where a stereotype is *mistaken*, we have to acknowledge what that stereotype *is*. Keeping things at a completely abstract level or, worse, pretending that there aren't any problems, is not an effective way to think about these issues.

It may help if we begin by noting that everyone—even white, heterosexual males—is to varying degrees the victim of stereotyping. At the university where I teach many of the students are from Oklahoma. People from other parts of the country sometimes categorize them as *Okies*, unsophisticated, uneducated, rural, maybe even ignorant and bigoted and dumb as dirt. This label is now sometimes used in a joking way, especially in Oklahoma, but if someone in Boston really has this stereotype of Oklahomans, it could affect whether a person from Tulsa gets a job they apply for or is admitted to the school they want to go to.

People in many parts of the world have negative stereotypes of Americans: we are very often seen as arrogant, insensitive, materialistic and selfish. In some places, e.g., some parts of the Middle East, the stereotype is even more negative: America is “The Great Satan”. Many of you are also stereotyped because of your ethnic background, physical appearance, religious views, and so on. There are also more specific stereotypes, like jock, ditzzy dumb blond, computer nerd, frat rat, and so on. All of these can lead others to treat you unfairly, on the basis of your perceived membership in a group, rather than on your own merits.

Stereotypes in many segments of our society have not changed dramatically over the last few decades, although many of them have become somewhat less negative. Not everyone has the same stereotypes for various groups, but many stereotypes are widely shared. Quickly write down the traits that you think are commonly ascribed to blacks or Hispanics or Jews or gays. Don't censor what you write (remember, you aren't saying that *you* feel this way about the group). You will probably come up with some of the following characteristics, which research shows are widespread aspects of common stereotypes. Blacks are stereotyped as athletic, hostile, and lazy. Jews as shrewd, intelligent, shrewd, and materialistic. Gay men as effeminate; lesbians as masculine.

## 25.4 Discrimination

When prejudice is widespread, it can lead to widespread, even institutionalized, discrimination. Less than fifty years ago in some parts of the United States Afro-

Americans were required by law to sit at the back of the bus, use separate water fountains, and attend separate schools from those used by white students. And they were treated this way *simply* on the basis of their membership in this group.

### 25.4.1 Overt vs. Modern Racism

**Modern racism:** less blatant racism that takes more subtle and indirect forms

In the last few decades such overt racism has decreased in many (though certainly not all) segments of American society. People have learned that such behavior is not viewed as socially acceptable, and they don't want to be labeled *racists*. Sometimes old racist attitudes remain, though now unexpressed, but many people's attitudes have genuinely softened. Even in places where blatant discrimination no longer exists, however, prejudice often persists and reveals itself in more subtle and indirect ways.

For example, people who would never dream of calling Afro-Americans 'niggers' may still assume that blacks can't do as well in school as whites. Again, probably very few of the doctors in the study mentioned at the beginning of this chapter were blatant racists; most probably thought of themselves as tolerant. Still, they behaved in ways that would have had much better consequences for whites than for blacks.

There is also less overt sexism in some segments of society, but modern (less blatant, more subtle) sexism remains. Modern racism and sexism are particularly dangerous, because it is harder for those who aren't the targets to see them. It is easier for them to argue that racism and sexism no longer exist.

### 25.4.2 In-groups vs. Out-groups

An *in-group* ("us") is a group that a person belongs to (or wants to belong to) and identifies with. We are all members of a variety of in-groups, from small social clubs or cliques to being citizens of an entire country. *Out-groups* ("them") are groups we don't belong to and don't identify with. The less we identify and sympathize with a group, or the more we dislike it, the more "out" it is. What's in and out depends on who you are. One person's in-group (e.g., the Ku Klux Klan) will be many others people's out-group.

There are various reasons why people identify with groups, including the fact that it helps foster and maintain self-esteem and a social identity. But whatever the exact causes, people think about their in-groups and out-groups in different ways.

Discrimination is often based on long-smoldering resentment that go back many centuries. Race is the most obvious, though not the only, example in the United States. Elsewhere it is religion: in Northern Ireland it is Protestants vs. Catholics; in India Hindus vs. Muslims. In India there are also divisions between



castes. In Africa there are tribal hatreds like those in Rwanda between the Tutsis and Hutus. And gender and sexual orientation evoke discrimination almost everywhere. In Asia, for example there are *76 million missing women*—76 million fewer women than there should be given the very nearly equal birth rates of males and females. This results from the fact that women in many societies are valued less than males, and so they receive less food, worse health care, and fewer resources generally.

### **Discrimination can be Easy**

Studies show that discrimination against out-groups increases when times are bad, when money or jobs or food become scarce. But it is also sobering to realize that prejudice and discrimination can also spring up quickly, almost spontaneously, in a wide variety of situations.

In Chapter 23 we saw how quickly the guards in the Stanford Prison Study came to see their prisoners as almost subhuman. We also saw that even young children are not immune. Remember Jane Elliott's third-grade class in Riceville Iowa (p. 476)? One day Elliott informed her students that brown-eyed children were smarter and better than blue-eyed children and so they should be treated better. The superior, brown-eyed students were accordingly given various privileges and the inferior blue-eyed students were subjected to rules that called attention to their new, inferior status.

Before the end of the day the brown-eyed students were banding together and discriminating against their blue-eyed former friends; they fought with them and ostracized them. Meanwhile the blue-eyed students became angry, sullen, and withdrawn. The next day the roles were reversed.

Years later many of the students felt that the experience had taught them things about discrimination that they would never have learned otherwise. If something like this can happen in a single day among friends, one can only imagine the effects of being treated like the brown-eyed children year after year in a setting that is never—unlike Elliott's experiment—"called off."

### **The Minimal Intergroup Discrimination Effect**

Many studies show that we have a strong tendency to identify with groups even when they are arbitrarily formed and short-lived. Groups formed on the basis of meaningless or trivial criteria are called *minimal groups*, and this identification is known as the *minimal intergroup discrimination effect*. Once someone becomes a member of a group, she will tend to favor that group even when the group is formed arbitrarily, as when students are randomly assigned, on the basis of coin flips that

they all observe, to one of two groups.

If a group member is then given a chance to divide rewards between members of her own group and the other group, most show marked in-group favoritism. This occurs even when she does not know the other members of her own group, will not get any of the rewards she is distributing to them, and her fellow group members cannot reward her in return. The mere membership in a group is often enough, without any selfish motives or long-term identification, to produce in-group bias. It is enough to make us like members of our group more and to treat them better.

## 25.5 Features of Stereotypes

We noted three key features of stereotypes above. They are (i) typically inaccurate, (ii) attribute the same features to virtually all members of a group (even though there are enormous differences among the group members), and (iii) are often very resistant to change. Since different stereotypes are inaccurate in different ways, issues of accuracy are best discussed when examining specific stereotypes. But we can say something general about the other two points.

### 25.5.1 Homogeneity

#### The Out-Group Homogeneity Bias

**Out-group homogeneity bias:** tendency to see out-groups as more homogeneous than they really are

The **out-group homogeneity bias** is the common tendency to see other groups (especially “out-groups”; e.g., people of other races, religions, countries, sexual-orientation) as *more homogeneous* (similar to one another) than they really are. Group members tend to see their own group as quite varied, whereas members of other groups are thought to be much more similar to one another.

Because the members of out-groups seen as more similar than they really are, it is easier to think that most of them fit a simple stereotype that lumps all of them together. We tend to think of this person as “just one more of them”—just another person from India, rather than as the unique individual Rajeev Singh.

Various things can account for the out-group homogeneity bias. In some cases people have little interaction with members of the out-group, and so they rely on a common stereotype that presents a simplified picture of the group. People may also be exposed to an unrepresentative sample of group members, perhaps because they only interact with a few people in the group or because they rely on media coverage that tends to focus on a small segment of it (e.g., athletes, executives, criminals).

### Strong Identification with One's Group

People's identification with an in-group can be very strong. *Ethnocentrism* is the belief that one's in-group (e.g., the people in one's country, region, religion) are *superior* to members of other groups. It need not involve ill treatment of those in other groups, but it often involves a distrust or wariness of them.

**Ethnocentrism:** belief that one's in-group is superior to other groups

*Communalism* is an extreme form of ethnocentrism. It is the tendency for people to divide into groups on the basis of religion, race, language, ethnic identity, or some other common feature. Insofar as it involves a pride in one's heritage and culture, there is nothing wrong with it, but it often spills over into prejudice and even violence against members of other groups. The "ethnic cleansing" in Kosovo or strife between Tutsis and Hutus in Rwanda or the closed aspects of various societies in the Middle East involve communalist tendencies. It is an irony of the modern world that increasing globalization and increasing communalism are both on the rise (in some cases the latter is surely a reaction to the former).

#### 25.5.2 Resistance to Change

Stereotypes and prejudices are often very difficult to change. It is frequently thought that many prejudiced people harbor prejudices partly because they have never really known members of the target group. Hence, if they can meet such people and interact with them in a positive setting, this will help them to see the group members as unique people, ones with many of the same features they themselves have, rather than simply as members of a group. This is often called the *contact hypothesis*: contact between individual members of groups in a positive setting where they are treated as equals will help each person become less prejudiced about the other.

Sometimes this works, but it is not as effective as we would hope. One reason for this is that prejudices often have a deeply-rooted emotion component, sometimes one going all the way back to a person's socialization as a young child. As a result of watching and talking to our parents and (later) others around us, especially our peer group, we internalize norms for thinking about different groups in different ways. But people also have cognitive (reasoning) biases and employ cognitive heuristics that help them retain prejudices even when they interact with members of the target group who don't remotely fit their stereotype. We will examine some of these in a moment, but first it will be useful to look at some of the cognitive mechanisms involved with prejudice.

## 25.6 Cognitive Mechanisms

### Different Types of Processing

There is increasing evidence that the human mind consists of a number of subsystems that work relatively independently of one another. For example, there seems to be a special system for recognizing faces that is largely independent of the perceptual system for recognizing other sorts of objects (one piece of evidence for this is that people with certain sorts of brain damage lose one ability while the other remains intact). Subsystems are not just somewhat independent; they can sometimes even arrive at rather different conclusions about the same thing.

Although the general topic of cognitive (mental) subsystems is fascinating, only a couple of views about alternative types of cognitive processes are relevant here. Patricia Devine has argued that our thinking about other people goes on in two different ways.

First, there is a relatively automatic, habitual mode of thinking that reaches conclusions quickly, involves little or no conscious thought, and over which we have little control. This system is useful, because we often need to draw conclusions fast; we don't have the time or attentional resources to reason in painful detail about everything we encounter, and this system lets us do a lot of social navigation on automatic pilot. But since it lies beyond our control, it can arrive at conclusions we would not endorse if we thought about things carefully; and since it lies outside our conscious thought, it can influence us without our knowing it.

Second, there is a more self-conscious mode of think that is more reflective and that we can control; it's what we normally think of as careful, conscious thought. It can lead to more well-reasoned conclusions than the automatic system provides.

### A Two-Stage Model of Information Processing

Devine puts these ideas together in a *two-stage model* of our processing of information about people. In the first stage, automatic processing draws quick conclusions. Here stereotypes can exert a strong influence. In the second stage, conscious processing can tell us that a conclusion we came to automatically is unsupported or even wrong, and we can then modify or ignore it.

Suppose, for example, that Wilma isn't prejudiced against blacks. Still, she has grown up in a society where the stereotype of blacks as hostile is very common so she has been exposed to it over and over. Under various conditions (e.g., when she sees a black or hears about one), this stereotype may be triggered automatically; it influences her thought without her intending for it to. It's automatic. But since Wilma is not prejudiced, her controlled, conscious processing can step in and re-

mind her that the stereotype is unfair and inaccurate, and so she can avoid acting on it. By contrast, if Wilbur is prejudiced, this automatic, stereotypical reaction may be the end of the matter; he won't go on to think about things more carefully.

In cases where we don't have much time or motivation to think about something consciously, e.g., because we are busy thinking about something else, we may only go through the first stage of processing. Devine's research suggests that under these conditions even people who are relatively unprejudiced may still be influenced by common stereotypes. In such situations a stereotype is activated automatically, and people don't go on to make the corrections they normally would have if they had given the matter conscious thought. Since the automatic processing occurs outside of conscious control, however, they aren't aware of this.

### Evidence for the Two-Stage Model

In one study Devine used a test to divide white subjects into a high-prejudiced group and a low-prejudiced group. She then displayed words on a computer screen very quickly so that subjects were not consciously aware of them, but the words were up long enough that they did register unconsciously. One half of each group got words that are commonly associated with the stereotype of blacks ('jazz', 'hostile', 'Harlem'); the other half got completely neutral words. She then asked the subjects to read a story about a fictional character named 'Donald'.

The story is designed to be ambiguous enough that it is possible for readers to interpret Donald's actions in various ways. Race is not mentioned in the story. But even low-prejudiced subjects who had seen the words involved in the black stereotype interpreted Donald in a much more negative way than those who saw the neutral words. The words associated with the stereotype somehow activated negative feelings. This automatic processing, completely outside the conscious awareness and control, influenced how they interpreted Donald and his actions.

Although the details of Devine's model aren't fully accepted by everyone working in her field, a number of subsequent studies confirm the basic idea that cultural stereotypes can operate automatically to influence how we think, without our even being aware of it. One moral of this research is the importance of careful, conscious thought, since without it we may arrive at conclusions that are quite different from the ones we would want to draw.

#### 25.6.1 Levels of Generality

There also seem to be differences in processing that are based on the *levels of generality* of the things that we think about. The social psychologists Fiske and Neuberg argue that the *attitudes* we have concerning *general categories* or groups

may differ a good deal from the attitudes we have toward their *individual members*. Many people claim to dislike lawyers *in general*, but they like their own (*individual*) lawyer just fine. Lots of people express contempt for politicians, but they like those that they know about; for example, most incumbents are popular in their home districts or states and win when they run for reelection. This is relevant here, because prejudices are often *general*, directed against “*them*”—that’s just what *they* are like.

The psychologist Robert Abelson suggests that there are two different subsystems here, and that they communicate poorly enough that they often lead us to have inconsistent attitudes (e.g., disliking a group while liking its members). A classic study by LaPiere in 1934 illustrates the idea. LaPiere traveled around the United States with a young Chinese couple. They were accommodated in all but one of the two hundred restaurants and hotels where they stopped.

Later LaPiere wrote these places and asked if they would be willing to accept Chinese guests. Of those who answered, 90% said “No.” There are probably many reasons for this, but one may be that the respondents had negative views about the *general* category of Chinese, but they didn’t have any difficulty with these two Chinese *individuals*.

If something like this is right, it helps explain why a positive interaction with individual group members may do little to change a person’s general attitude toward the group (we will see additional explanations below). The reason is that the cognitive subsystem that deals with generalities and groups is to some degree unaffected by the subsystem that deals with specific individuals. If so, then it may be more effective to change prejudices against groups by changing how people think about the group’s history or achievements, rather than only by getting people to interact with group members; the latter is important, but it may not be enough because it may not affect the cognitive subsystem that is responsible for how we think about groups.

## 25.7 Flawed Reasoning and Prejudice

Many fallacies and cognitive biases help foster and maintain prejudices

Many of the cognitive biases, fallacies, and other pitfalls in reasoning studied in earlier chapters help foster and maintain stereotypes and prejudices. We will now examine some of them.

### Belief Perseveration

*Belief perseveration* (Ch. 8.5) is the common tendency to retain a belief even after our original reasons for thinking it true have been undermined. Some beliefs are

so thoroughly entrenched that they are impervious to evidence that would discredit them. Once we acquire a belief, it can be very difficult to get rid of it, even when the evidence clearly tells us that we should. The next two cognitive biases help explain why this is so.

### Confirmation Bias

Many studies (as well as a bit of careful observation) document our tendency to look for and remember positive evidence that supports our beliefs while not seeking (or even avoiding) negative evidence that could tell against them. *confirmation bias* (Ch. 18.5) is the selective tendency to look for, notice, and remember confirming or positive evidence (that supports what we think) while ignoring or downplaying disconfirming or negative evidence (which suggests that what we think is wrong). For example, if Wilbur already thinks that women are bad drivers, he is more likely to notice or remember cases where women drove badly than cases where they did not.

Confirmation bias leads to bad reasoning, because it means that we never put our views to the test. Hence there will be little chance of finding out that we are wrong, even when we are. This bears on the topics in this chapter because if we have stereotypes, we will tend to look for (or notice) evidence that confirms them, and we won't look for (or notice) evidence that might show they are mistaken. So even if we encounter group members who don't fit our stereotype (i.e., people who disconfirm it), we may overlook or ignore it.

### Subtyping

Sometimes when people interact with members of a stereotyped group it is so obvious that the group member doesn't fit the stereotype that it is impossible for them to ignore it. But even this often fails to lead to any change their stereotype. Instead they maintain their stereotype by appealing to subtypes.

A **subtype** is a stereotype for what is seen as a small and atypical subset of a group. For example, research shows that prejudiced whites have a stereotype of blacks as lazy and hostile. When someone like this meets a black who obviously doesn't fit the stereotype—say someone who worked their way up from a life of poverty to be CEO of a large, successful corporation—the stereotyper often classifies the other person as a member of a *subtype*, like Black Businessman. This allows her to maintain her stereotype in the face of contradictory information (“this person doesn't fit the profile, but he's an exception—most blacks do”). There are also gender subtypes like career woman, feminist, and housewife or jock, playboy, and punk.

**Subtype:** a stereotype for what is seen as a small and atypical subset of a group



### Sample Size

When we draw inferences from small or unrepresentative samples, our conclusions will be unreliable (Ch.15.2.3). Some stereotypes no doubt arise from such hasty or unwarranted generalizations. In any large group, there are bound to be some members with negative characteristics. If we happen to interact with just these members of a group, it is easy (though unwarranted) to generalize to the entire group from this faulty sample.

### Heuristics

The uses of various cognitive heuristics (like availability and representativeness) (Ch. 17) have associated patterns of faulty reasoning called biases, and these can also lead to bad reasoning about other people.

### Two Meanings of ‘Bias’

Before seeing how this works, it is important to be clear that we are dealing with two different meanings of the word ‘bias’ here. When we studied heuristics we also examined their associated biases; for example, the representativeness heuristic often leads to a bias to think that a conjunction can be more probable than either of its conjuncts. In this context we often speak of a ‘cognitive bias’, and ‘bias’ means a systematic tendency to reason badly in a certain way; for example, we have a bias, a tendency, to think conjunctions can be more probable than their conjuncts and we have a tendency to overlook regression to the mean.

By contrast, when we discuss prejudices, we are likely to talk about ‘biases’ towards groups. Here the word ‘bias’ means pretty much the same thing as ‘prejudice’. Context will make it clear which sense of the word is involved.

### Availability

We use the availability heuristic (Ch. 17.2) when we base judgments of frequency on the ease with which examples can be recalled or imagined. This heuristic inclines us to assume that things that are easier to remember or imagine are more likely to occur.

Often we remember certain things because they really do occur frequently, and when this is the case, the available sample will be a good enough for the rough-and-ready inferences of everyday life. But things may be available in memory or imagination for reasons having little to do with their true frequency or probability. In such cases the availability heuristic will lead us to rely on a sample (the cases we easily remember) that is unrepresentative in various ways.

#### Two meanings of ‘bias’

1. tendency to employ a certain pattern of bad reasoning
2. prejudice



If we expect members of a group to behave in a certain way (as we will if we are strongly influenced by a stereotype of the group), we will be more likely to notice and remember behavior that matches our expectations. If Wilbur thinks that women are bad drivers, he will be more likely to notice when women do drive badly than when they don't (or than when men drive badly), and he is more likely to remember those occasions where there do. Such cases will be more *available* to Wilbur when he thinks about women drivers, and this will lead him to think that a larger percentage of women are bad drivers than is actually the case.

Availability may also explain some cases of the out-group homogeneity bias. If we only interact with a small subgroup of a group, or if we rely on media reports that deal only with a small segment of the group, then our sample will be small and unrepresentative. It will easily come to mind, however, and so we may conclude that most members of the group are like the members of the small (and quite possibly atypical) subgroup that we know about.

### Representativeness

We use the *representativeness heuristic* (Ch. 17.3) when we conclude that the more like a representative or typical member of a category something is, the more likely it is to be a member of that category. Put in slightly different words, the likelihood that  $x$  is an  $A$  depends on the degree to which  $x$  resembles your typical  $A$ . So here we reason like this:  $x$  seems a lot like your typical  $A$ ; therefore  $x$  probably is an  $A$ . For example, Linda *resembles* your typical feminist (or at least a *stereotype* of a typical feminist), so (many of us conclude) she is likely to *be* a feminist. Sometimes this pattern of inference works, but as we have seen, it can also lead to very bad reasoning.

This can lead us to believe that someone is a member of a group when they are not ("they fit the group stereotype, so they probably belong to the group"). We also frequently make the opposite mistake. Someone is in the group so, we (mistakenly) conclude, they probably fit the stereotype.

### Anchoring and Adjustment

An anchor (Ch. 17.5) is a sort of reference point. Often we know that an anchor is too high or too low, and we try to adjust for it. But very often we *don't adjust enough*. So although we don't simply accept the anchor as given, it still skews our final judgment; a high anchor often leads us to overestimate something, and a low anchor leads us to underestimate it.

If we frequently hear that members of a particular group are unusually dumb or dishonest we may assume that this claim is too extreme. But it may still provide

an anchor, so when we adjust away from it we don't adjust enough. We don't think the group is as dishonest as we are told, but we may still think that they may be a bit more dishonest than average—even when they aren't. The anchor biases our final judgment.

### Memory Effects

Earlier (p. 131) you read about people who heard a fairly neutral description of the early life of Betty K. The story contained sentences like 'Although she never had a steady boyfriend in high school, she did go on dates'. Later, half of the subjects were told that Betty later became a lesbian while the other half were told that she got married. The first group was more likely to remember "She never had a steady boyfriend" rather than "she did go on dates." In the second group the results were reversed. What the people in the first group *remembered* was influenced by their *stereotype* of a lesbian.

Memory is an active, constructive process. Among other things We fill in gaps in memory in ways that help us make sense of the world. For example, people heard a fictitious story about a dictator. In one version he was called 'Gerald Martin'; in another he was called 'Adolph Hitler'. The story didn't mention Jews, but many of the people who heard the Hitler version of the story later thought—"remembered"—that it contained the sentence 'He hated the Jews and persecuted them'. People in the other group did not. People in the first group filled in details based on their knowledge of Hitler. We can also fill in details about a person based on our beliefs about the groups they belong to. Here our stereotypes affect the way we edit memories and fill in details.

### Collective Memory

Much as individuals have memories that are stored in their brains, societies have "collective memories" (Ch. 8.10) that are embodied in their beliefs, legends, and stories about the past. Social scientists have found that collective memories change over time and that they are often quite different from the original events that gave rise to them.

Sometimes people in power, particularly in totalitarian societies, set out to revise collective memory. There are many techniques for doing this, including rewriting textbooks, constantly repeating the rewritten version of history, and forbidding discussion of what actually happened. But it also occurs, in a less conscious way, in more open societies. It often happens that certain groups (e.g., Native Americans) are portrayed in inaccurate ways in a group's collective memory. This can certainly affect how people will think about that group, what images of it will be

available, and how they will feel about it.

### **Illusory Correlation**

We believe in an *illusory correlation* (Ch. 15.4) when we mistakenly think that two things go together when in fact they are unrelated. For example, Wilbur may think that being a woman and being a bad driver tend to go together, i.e., that these two things are positively correlated. Illusory correlations often result from an overemphasis on positive cases, and Wilbur is likely to notice the cases where the two things go together (cases where woman drive badly) and to ignore cases where they don't (cases where women drive well, or cases involving men's driving). He ignores evidence that tells against our view that the two things are related. It is often easier to think of positive cases where two things go together than to think of negative cases where they don't, especially if we already think that two things typically accompany each other

There are many reasons why people hold stereotypes, but belief in illusory correlations often reinforces them. Thus people may believe that members of some race or ethnic group tend to have some characteristic—usually some negative characteristic like being lazy or dishonest—which is just to say that they believe that there is a (positive) *correlation* between race and that characteristic. We are especially susceptible to a belief in an illusory correlation when it involves things that are easily noticed, and members of many (though not all) minorities are.

### **The Validity Effect**

Researchers have found that the mere repetition of a claim will lead many who hear it to think that it is more likely to be true than they would have supposed if they hadn't heard it before. This is called the *validity effect* (Ch. 18.7). The mere repetition makes the claim seem more likely to be true (or "valid"). This effect occurs with true statements, false statements, and statements that involve expressions of attitudes.

Since the validity effect often leads people to believe things without giving them any thought, if we hear a claim over and over, from the time we are small children, we will have some tendency to believe it simply because we have heard it so much. So if people repeatedly hear that jews are materialistic, without hearing others say that they aren't, they will have a tendency to believe it.

### **The Just-World Hypothesis**

Many people think that the world is basically fair and just. People usually get pretty much what they deserve and deserve what they get. The psychologist Melvin

Lerner called this the *just-world hypothesis* (Ch. 18.8).

Lerner has shown that when people learn about an unfair outcome that is otherwise difficult to explain, they look for a way to blame the person who suffers the misfortune (“they must have done *something* to bring on their problems”). To the extent that we feel this way, we will tend to think that most people who aren’t doing well are getting what they have coming (until fairly recently many rape victims were viewed in this way). So if a group is treated badly, we may feel, they must have some defects that explain this bad treatment (“Where there’s smoke there’s fire”). In short, we tend to blame the victim, and sometimes this extends to blaming entire groups.

### **Wishful Thinking**

We engage in *wishful thinking* (p. 163) when we disregard the evidence and allow our *desire* that something be true to convince us that it really *is* true. Our strong human tendency to wishful thinking is one reason why claims by pseudoscientists, advertisers, and others are accepted even when there is little evidence in their favor.

When things go badly *for others* we tend to blame the victim. But when things go badly *for us* we often look for *someone else* to blame. We would rather not think that we are at fault, or even that we had a run of plain bad luck. **Scapegoating** is blaming others for problems, especially social or personal problems. Prejudice, stereotypes, and scapegoating reinforce each other. Unflattering stereotypes about a group lead to prejudice against its members, and the need to find scapegoats provides an emotional cause for adopting stereotypes.

Prejudices and scapegoating seem to be more common when times are tough, for example when the economy is bad. A classic example is the way that Nazis in Germany in the 1930s and 40s blamed the Jews for many of their social problems (which included a devastated economy). Scapegoating is still very common today, and with the rise of communalism it may increase.

### **Dissonance Reduction**

Most people do not see themselves as evil or sadistic. Consider people who find themselves in a situation where they systematically treat others badly, or even where they simply ignore their plight. It would lead to a good deal of cognitive dissonance (Ch. 19) if they thought to themselves: although I am basically a good person, I am treating these people badly (or ignoring their suffering), and they do not deserve to suffer. One way to reduce this dissonance is to change one’s attitude toward the mistreated group: “Maybe they really do deserve to be in their situation after all”; “Maybe (in extreme cases) they aren’t even quite human.”

This fits with the common finding that most of the people who worked in Nazi concentrations camps came to see their victims as less than human. When you treat someone badly, there is a tendency to derogate them, to think “well, they deserved it.” When we are the ones victimizing someone, blaming the victim helps us reduce dissonance.

### **Inert Knowledge**

As with any knowledge, we may know that common stereotypes are inaccurate or oversimplified in various ways, yet fail to keep this in mind outside the setting (e.g., a classroom) where we learned it. When we need it, the knowledge just doesn’t come to mind. We aren’t aware of our situation as one where our knowledge about the limitations of a stereotype is relevant; we simply don’t think about such things.

### **Framing Effects**

We have seen more than once that the way things are presented or worded can influence how we think about them. A description of a welfare program as one that helps young children get a hot breakfast will evoke very different images from a description of it as an inducement for Welfare Queens to have more children so they will receive bigger checks. In highly emotional issues of the sort that surround race, sexism, sexual orientation and the like, different people will frame things in vastly different ways.

We have also seen that it is extremely important whether something is framed as a gain (e.g., lives saved) or as a loss (lives lost). Most things can easily be framed either way. For example, a member of a dominant group can frame something as a higher employment rate among a minority group (a gain) *or* as members of the minority group taking jobs away from members of the dominant group (a loss).

### **Status Quo Bias**

The *status quo* bias is a preference to keep things pretty much the way they are. Unless things are going badly, people often prefer to keep things the same rather than risk trying something new. So people in groups that are doing well will have a preference to keep things largely as they are, rather than to risk new policies (e.g., a new affirmative action program) that might make things worse for them.

### **Contrast Effect**

The contrast effect occurs when our evaluations of, or judgments about, something are influenced by the contrast between it and the things around it. If I view another

group (and its members) as inferior to my group (and its members, including me), then I look better by comparison. It may increase my sense of self-worth and self-esteem. So people sometimes derogate other groups so that their own group will, by contrast, look better.

### **Self-fulfilling Prophecy**

A self-fulfilling prophecy is the tendency for a person's expectations about the future to influence that future in a way that makes the expectations come true. Sometimes we have expectations about people that lead us, unwittingly, to treat them in a certain way, and treating them in this way may then lead them to behave in the way that we thought they would.

For example, if I have a stereotype of Afro-Americans as hostile, I may be inclined to be hostile to John just because he is an Afro-American. And this may lead him to react with hostility, even though he would have been friendly if I'd been friendly myself. My prediction leads me to act in a way that makes my prediction come true.

Rosenthal and Jacobson's work (p. 361) is relevant here. Recall that they told grade school teachers at the beginning of the school year that their incoming students had just been given a battery of tests and that twenty percent of the students should be expected to blossom academically in the coming year. In fact, the students in this group were selected at random.

Nevertheless, these twenty percent ended up improving more than the other students. Teachers expected the students in the targeted group to blossom, which led them to act in ways that encouraged the students to do so. This sort of self-fulfilling prophecy is sometimes called the *pygmalion effect*, and many studies since have documented its existence.

### **Stereotypes as Self-fulfilling Prophecies**

Stereotypes can also serve as self-fulfilling prophecies. If teachers expect students from some groups to perform better than others, this may lead them to treat their students in ways that will make these expectations come true. In a society where people think that women are incapable of careers working with computers, young girls are likely to be treated in a way that suggests they can't do such work. Furthermore, the interest they display in computers may be discouraged, and they will be encouraged to adopt different interests. Years of such treatment *will* make it much more difficult for a woman to have a career working with computers.

In studies by Word, Zanna and Cooper, it was shown that white interviewers conducting mock (practice) job interviews were more at ease with white inter-

viewees than with black ones, and this led the white interviewees to be more at ease while the black interviewees were less so. The interviewers' expectations that white interviewees would be easier to talk to led them to behave in ways that tended to make this expectation come true.

Self-fulfilling prophecies can lead us to behave in ways that lead others to confirm a stereotype. Then, because of confirmation bias, we may focus on this evidence, without thinking to look for evidence that would *disconfirm* our stereotype, and without asking how others might have behaved *if* we had behaved differently toward them.

### The Power of the Situation

We have seen (Ch. 23) the very strong power of social situations to influence how we think about people and how we explain why they do the things they do.

### Conformity Pressures

As we grow up we acquire many beliefs and attitudes and norms without ever thinking about them. This process is called *socialization*. Without it we would never become full members of a culture, which is to say we would never become fully human. But we can be socialized to hold many different beliefs, and if someone grows up always hearing that women are inferior to men (and almost never hearing that they aren't), they have to be an exceptional person not to be influenced by this.

### Situational vs. Dispositional Explanations

In Chapter 23 we saw the importance of the distinction between internal, dispositional explanations of the causes of someone's actions, on the one hand, and external, situational explanations, on the other. Dispositional explanations cite internal causes like character traits, whereas situational explanations cite features of the circumstances surrounding the actor.

**Internal Causes:** Causes "inside" the person: his or her personality traits or characteristics or dispositions and his attitudes, values, desires.

1. John did poorly on the exam because he is *stupid* and *lazy*.
2. Wilma screamed at Wilbur because she is a very *hostile* person.
3. Wilbur paid Edna that compliment because he is *generous* and *kind*.

**External Causes:** Causes "outside" the person: features of the *situation* or circumstances in which the person acts.



1. Wilma didn't help the victim because the *situation* was one in which it was unclear whether the victim needed help and no one else at the scene seemed to think that he did.
2. John said that line 2 matched line A because of the strong *social pressure* exerted by the other people in the experiment.
3. Wilbur followed the leader's orders, but *anyone else* would have done the same thing in those *circumstances*.

All actions take place in some context or situation, and *both* a person's internal states (dispositions, attitudes, etc.) *and* features of the situation (e.g., the presence of others, the commands of an authority figure) always play some role in determining what he does in that situation. It is *never* the case that internal causes aren't important. But people strongly *overestimate* the importance and strength of internal causes while underestimating the importance and strength of external causes.

This is called the *fundamental attribution error* (p. 477). We commonly underestimate the power of the situation to influence behavior. The two types of explanation are also relevant to the *actor-observer asymmetry* (Ch. 23) We tend to see other people's behavior as internally caused, but to see our own as externally caused.

The distinction between internal (dispositional) explanations and external (situational) explanations is also important when we think about members of groups. For example, the average score for various minority groups on certain standardized tests (like the ACT) is lower than the average score for whites. But what does this mean? The issue is one about whether a dispositional explanation (members of the lower-scoring groups just aren't as smart) or a situational explanation (members of the lower-scoring group had to attend much poorer schools) is the right one.

More generally, explanations that cite stereotypes cite internal, dispositional causes; a group member does certain things because that it just the sort of person that he (and the others in the group) *are*. *They* are just hostile or lazy or materialistic or bigoted: "He did that because he's an Okie, and all Okies are bigoted".

Blaming the victim also involves giving an internal, dispositional explanation of the victim's plight. They suffer their misfortune because of the sorts of people they are ("O.k., she didn't *deserve* to be raped, but anyone who behaves like she did is asking for trouble"). Whereas in many (though of course not all) cases people suffer misfortunes because of features of their circumstances that are beyond their control, and here external, situational explanations will be more accurate ("She just happened to be there when the prison escapee came along").



### The Ultimate Attribution Error

Thomas Pettigrew has called our tendency to give dispositional explanations for the negative or unsuccessful behavior of an *entire* group of people the *ultimate attribution error*. By contrast, although we often give dispositional explanations to explain the failures of the members of an out-group, we tend to give external, situational explanations of their positive behavior and successes (“He was just lucky,” “She must have gotten some special break because she was a woman”).

**Ultimate attribution error:** giving dispositional explanations for the behavior of an entire group

There is evidence that when a member of a stereotyped group behaves in a way that *disconfirms* a negative stereotype, we tend to offer a *situational* explanation of her behavior; we reason: “Almost anyone would do the same thing in that situation, and so it doesn’t really tell us much about the person involved.” This is yet another way in which we can retain our stereotype in the face of disconfirming evidence. For example, Wilbur is convinced that women are bad drivers. One day he sees Sue think quickly, then steer out of the way of an on-coming truck. But, he reasons (situationally), “anyone in her position would be so frightened that they would be really focused and alert, and so they’d do the same thing.”

## 25.8 Responses

### 25.8.1 Remedies and Reasoning

Problems stemming from prejudices and discrimination are among the most severe difficulties we face today. There are no easy answers or quick fixes, but it is important to look for strategies that could help reduce these problems.

Although mere contact between members of different groups is not as effective at changing stereotypes or reducing prejudice as we might have thought, there is something to the contact hypothesis. When members of groups that dislike or distrust each other interact in a positive setting, it often does help members of each group see members of the other group in a more positive light. But if people are brought together and the setting isn’t positive, things may actually get worse, and it isn’t always easy to ensure a positive setting in the real world.

For example, contact between blacks and white increased when schools were desegregated, and this did sometimes lead to less prejudice. The general findings suggest that this contact did not significantly reduce prejudice, however, and in some cases prejudice even increased. One reason for this may be that the contact in such settings is often fairly minimal; there is still a lot of segregation in desegregated schools. Blacks tend to congregate with blacks, Hispanics with Hispanics whites with whites (a quick look at who sits with whom in most school cafeterias should convince you of this). But even so, contact did increase without any

noticeable drop in prejudice. So mere contact is not enough.

### **Education and the Jig-saw Classroom**

Education is probably the best hope for changing stereotypes and prejudices. In general younger people in our culture have less negative stereotypes than older people, and it is probably easier to change their attitudes than it is to change the attitudes of people who are more set in their ways.

One setting that has had a good deal of success is the *jig-saw classroom*. It was developed by Elliot Aronson and his co-workers in the 1970s in Austin, Texas. The idea is to set up a situation where two or more groups of students must work together to achieve a common goal. Students are placed in small groups. Each person in the group has to master a given bit of the material and teach it to the other people in the group, so each person and their contribution is like a piece of a jigsaw puzzle that must be worked in to complete the picture.

For example, in a project on the life of George W. Bush one student might gather information on Bush's childhood, another on his time in college, and so on. So the success of the group requires the cooperation of all of its members. Such approaches are now part of what has become known as the cooperative learning movement.

This interdependence requires working together, and it often does foster mutual interest and respect that are not present in more traditional classrooms. Even approaches like this are not a universal cure, however, since the quicker students in a group may actually come to resent group members that they see are slowing the group's work down.

### **Breaking Habits**

Stereotypes and prejudices are *habitual*, relatively automatic patterns of thinking and feeling, and Patricia Devine has usefully compared changing a stereotype to breaking a habit. Often stereotypes are very ingrained, and we can't expect to get rid of them easily—any more than a long-time smoker can hope to easily kick the habit. It takes work over a period of time, and even then the results may be less than perfect. But if one wants to change how they think about another group, conscious effort to do so, including catching oneself when one falls back into stereotypical modes of thought, can help. Of course none of this works unless a person wants to change.

## Conclusion

In this chapter we have seen that stereotypes (and through them prejudices) are usually based on flawed or fallacious reasoning. Thinking more clearly about them and the evidence on which they rest is the first step in dealing with them.<sup>1</sup>

## 25.9 Chapter Exercises

1. What is the out-group homogeneity bias? Explain some cognitive (i.e., reasoning) mechanisms that could lead to it. Then explain some ways in which it could contribute to prejudice.
2. Give an example where the use of the availability heuristic helps foster a stereotype and explain how it does so. How could the danger of this be minimized?
3. What is a genuine correlation? What is an illusory correlation (and how does it differ from a genuine one)? Explain how such cognitive illusions could lead to prejudice and give a concrete example.
4. Many inner cities have become so violent that the leading cause of death of young black men there is now murder. Give an example of an internal explanation of this that involves a common stereotype of black males. Then give an external, situational explanation of the same phenomenon. How do the two explanations differ? What is their relationship to the ultimate attribution error?
5. List two stereotypes that are commonly associated with two groups. In what ways are they unfair? Identify faulty types of reasoning that may have led to the formation of each of these stereotypes, or that may have helped keep them in place once they were present.

---

<sup>1</sup>For a thorough discussion of the estimates of the size of Native American populations at various times see Russell Thornton, *American Indian Holocaust and Survival*, University of Oklahoma Press 1987. Robert P. Abelson's "A Personal Perspective on Social Cognition," in Patricia Devine, ed., *Social Cognition: Impact of Social Psychology*, Academic Press, 1994; 15-37, contains an excellent introduction to cognitive subsystems and their possible implications for prejudice and discrimination. Any good social psychology text will discuss many of the psychological aspects of stereotypes and prejudices. One that is especially good on these topics is Elliot Aronson, Timothy D. Wilson & Robin M. Akert, *Social Psychology: The Heart and the Mind*, Harper Collins College Publishers, 1994. See Patricia G. Devine and Andrew J. Elliot, "Are Racial Stereotypes Really Fading? The Princeton Trilogy Revisited," *Personality and Social Psychology Bulletin*, 21, 1995; 1139-1150 for a sobering discussion of just how negative contemporary stereotypes of blacks still are. Further references to be supplied.

6. What is the difference between an internal (or dispositional) explanation of an action and a situational explanation of it? What roles do these types of explanation play in the fostering and the maintenance of prejudice?
7. Give an example (real or imaginary) where the use of the representative heuristic helps foster or maintain a prejudice or stereotype and explain how it does so. How could the danger of this happening be minimized?
8. What is the ultimate attribution error? Explain how it is similar to (and also how it differs from) self-serving biases, on the one hand, and how it is similar to (and also how it differs from) the fundamental attribution error, on the other. How can the ultimate attribution error lead to prejudice?
9. What is modern racism and how does it differ from more the more traditional sorts of racism prevalent in our country several decades ago? Give a concrete example of modern racism.
10. How might dissonance reduction lead to prejudice? Explain the general mechanisms and illustrate your discussion with a concrete example.
11. What is the jigsaw classroom? How might the ideas it embodies be applied in other settings? Explain and give a concrete example.
12. Give an example of scapegoating that has occurred in recent years. What part did stereotypes play in it? In what ways did it involve faulty reasoning?
13. Give an example of communalism and note any harmful aspects of it.
14. Analyze each of the following dialogues.

*Edna:* I hear you just got hired as the manager of that new pizza place on Lindsey.

*Wilbur:* Yeah, and it's a fantastic job. The only downside is the boss says I have to hire some women delivery drivers.

*Edna:* So what's wrong with that?

*Wilbur:* Well, don't get me wrong. I *like* women. But you know what kind of drivers they are. And they'll spend twenty minutes talking to the people at each stop. They'll be lucky to make two deliveries an hour. I mean, just look at my sister Sue. Four wrecks in just two years. And she can't stop talking.

*Sue:* What's up Edna? You look steamed.

*Edna:* Wilbur just drives me up the wall. He's so prejudiced. He's just like all men, though, such a sexist.

*Professor Smith:* Those right wing fundamentalists are really getting out of hand. Of course I support their right to free speech, but why do they have to use it to say such idiotic things? They must have IQs of about 40.

*Professor Jones:* I don't like their views a lot myself, but some of them are extremely intelligent. That new student from Louisiana is a fundamentalist, and he's the smartest person we've ever had in our graduate program.

*Professor Smith:* Well, if he's a fundamentalist, maybe you'd better reevaluate him.

*Bubba:* Thank God somebody finally stood up for what's right.

*Sue:* What do you mean?

*Bubba:* That bar down the street throwing out the gays and lesbians. They should all be shot.

*Sue:* What have they ever done to you?

*Bubba:* Well, I've been lucky enough not to know any, but I know all about them. It's like a disease. It's one of the reasons the economy is going to hell in a hand basket.

*Sue:* But Tom's gay, and everybody likes him.

*Bubba:* Well, I don't mean him. He's different. But you know what I mean.

15. Write a dialogue that illustrates the use of a stereotype, a second that illustrates scapegoating, and a third that illustrates a self-fulfilling prophecy.

16. Discuss some ways that more careful attention to evidence and reasoning might help reduce prejudices.



# Chapter 26

## Social Dilemmas

**Overview:** Social Dilemmas are situations in which actions that seem to be in each individual's own self-interest lead to outcomes that are worse for everyone. Such situations occur in international relations, in countries and cities, in families, and in two-person interactions. Their common structure is clearly displayed by a game called the *Prisoners' Dilemma*, so we will begin with it. We will then examine social dilemmas in a variety of settings. We conclude with a discussion of strategies for extricating ourselves from such dilemmas.

### Contents

---

<b>26.1 Prisoner's Dilemmas</b> . . . . .	<b>522</b>
<b>26.2 Real- Life Prisoner's Dilemmas</b> . . . . .	<b>524</b>
26.2.1 Public Goods and Free Riders . . . . .	525
26.2.2 The Assurance Problem . . . . .	527
<b>26.3 How is Cooperation Possible?</b> . . . . .	<b>530</b>
26.3.1 Coercion . . . . .	531
26.3.2 Positive By-Products of Cooperation . . . . .	532
26.3.3 Prudence and the Prospect of Future Interaction . . . . .	533
26.3.4 Loyalty . . . . .	534
26.3.5 Moral Principles and Individual Ideals . . . . .	534
<b>26.4 Fundamental Values: Clashes and Tradeoffs</b> . . . . .	<b>535</b>
<b>26.5 Chapter Exercises</b> . . . . .	<b>538</b>

---

## 26.1 Prisoner's Dilemmas

The notorious felons *A* and *B* rob a bank. A few days later they are apprehended and booked into jail. The DA feels certain that they committed the robbery, but she doesn't have enough evidence to convict either prisoner unless one of them confesses. She puts the two prisoners in separate interrogation rooms and makes the following offer to each.

1. If you confess but the other prisoner does not, you can go free (in return for turning state's evidence). The other prisoner will be sentenced to 10 years.
2. If you both confess, you each get 5 years.
3. If neither of you confess, you will both get 1 year (I can't get a conviction on robbery, but both of you were carrying unlicensed concealed weapons when you were arrested).

The DA informs each prisoner that she has made exactly the same offer to the other prisoner, then leaves him to ponder his options.

Each prisoner may feel some loyalty to the other prisoner or may consider the possibility that if he squeals the other will seek revenge. But suppose each is so terrified of prison that such considerations play a negligible role in his deliberations. What would you do?

A two-by-two matrix (Figure 26.1, p. 523) allows us to visualize the possible outcomes. The rows of the matrix represent *A*'s options (Confess, Don't Confess). The columns represent *B*'s options (the same as *A*'s). Each of the four cells of the matrix specifies one of the four possible outcomes. The first number in each cell represents the number of years *A* gets, and the second number in the cell represents the number of years *B* gets.

Thus the top left-hand cell of the matrix represents the outcome where *A* and *B* both confess. The numbers in this cell are 5, 5, which indicate that in this situation both prisoners get a 5 year sentence. The top right hand-hand cell represents the outcome where *A* confesses but *B* does not. The numbers in this cell are 0, 10, which indicates that in this situation *A* gets 0 years and *B* gets 10. The two bottom cells work the same way.

Now put yourself in *A*'s position. How would you reason? *A* doesn't know whether *B* will confess, but *B* only has two options: either he confesses or he does not. So *A* considers each possibility in turn.

1. ***B* confesses** (this means we are considering the first column). In this condition I get 5 years if I confess and 10 years if I do not. So if *B* does confess, I'm better off (by five years) if I confess too.
2. ***B* does *not* confess** (we are now in the second column). In this condition I



		<b>B</b>	
		Confess	Not Confess
<b>A</b>	Confess	5, 5	0, 10
	Not Confess	10, 0	1, 1

Figure 26.1: Prisoner's Dilemma

get 0 years if I confess and 1 year if I do not. So if *B* doesn't confess, I'm better off (by one year) if I confess.

Either way *A* is better off confessing. So confessing looks like the rational thing for him to do. *B* will go through exactly the same pattern of reasoning, so it's also rational for him to confess.

The unsettling result is that when each prisoner pursues his own-self interest (looking out for Number One), the outcome is collectively self-defeating. They both end up with a condition (5 years in prison) that is four year's worse than the fate they would have had (1 year) if they'd cooperated. Each prisoner acts to maximize his own self-interest, so from the point of view of each prisoner it makes good sense to confess. But the result for each is much worse.

Would it make a difference if the prisoners were allowed to communicate before making their decisions? It depends. In many cases people reach an agreement to do something at a later date. Suppose that is the case here. *A* and *B* discuss the DA's offer and agree to cooperate.

But back in his cell *A*'s thoughts return to his earlier line of reasoning. He asks himself: Is it rational for me to keep my agreement? And the same line of reasoning he went through before convinces him that he is better off breaking his promise to *B* and confessing. Since *B*'s reasoning will parallel *A*'s, both confess. The ability to communicate isn't enough to solve the problem; if *A* and *B* do not trust each other, then their most fervent promises of cooperation won't help.

## 26.2 Real- Life Prisoner's Dilemmas

There are many situations like the prisoner's dilemma in the real world, many real-life situations in which the pursuit of narrow self-interest leads to an outcome that is worse for everyone. Since this pattern is so clearly displayed in the Prisoner's Dilemma, such situations are often called Prisoner's Dilemmas as well.

### International Relations

For several decades the United States and the Soviet Union were locked in a nuclear arms race. This cost vast amounts of money and led to a situation where the complete annihilation of the world was a real possibility. Both countries would have preferred mutual disarmament to the spiraling arms race. But like our prisoners *A* and *B*, they were trapped in a dilemma.

Each reasoned the same way. Either the other side will keep its nuclear arsenal or it won't. If they keep their weapons, then we must keep ours (so we won't be at their mercy). On the other hand, if the other side disarms, it's still good for us to have nuclear warheads so that they will have to stay in line. A perfectly rational line of thought leads to an outcome that is worse for both parties. The situation is depicted in Figure 26.2 on the facing page; the figure below; here the top entry in each cell indicates the outcome for the U.S., and the bottom entry in the cell indicates the outcome for the U.S.S.R.).

The general pattern in a prisoner's dilemma can arise when more than two parties are involved. The above line of reasoning would be essentially the same if there were more nations, say the United States, the Soviet Union, and China, engaged in a nuclear arms race. All three countries would be better off if everyone disarmed, but the situation makes this almost impossible. Indeed, the basic structure of the situation would be the same for any number of parties.

Cartels face a similar problem. The countries in OPEC (an oil cartel) reach an agreement to limit their production of oil in order to keep the price reasonably high. But it is impossible to determine precisely how much oil any particular country sells, and so each country will have an incentive to cheat.

		USSR	
		Build Arms	Disarm
USA	Build Arms	Costly & dangerous, Costly & dangerous	Dominate the world, At the mercy of the other
	Disarm	At the mercy of the other, Dominate the world	Cheaper & world safer, Cheaper & world safer

Figure 26.2: The Arms Race

Why? Because each country can reason as follows. Either most of the other countries will stick to their production quota, or they will cheat. If most cooperate, the price of oil will remain high, and we are better off producing more to bring in extra money. On the other hand, if most of the other countries cheat, we are suckers if we don't; in this case we had better get as much oil on the market as fast as possible, before the price drops through the floor.

This leads to a situation where many of the countries are likely to cheat, with the result that the price of oil goes down and they all suffer. Each of the parties would have been better of sticking to their original agreement, but strong pressures to defect are inherent in the situation. Many other international situations, e.g., those involving trade negotiations and world standards for pollution reduction, generate similar difficulties.

### 26.2.1 Public Goods and Free Riders

Prisoner's Dilemmas are a danger in any situation involving public goods. The following four features are each necessary (and together jointly sufficient) for some-

thing's being a **public good**:

1. Action by many of the members of a group is required to provide that good; the action of just a few members isn't enough.
2. If the good is provided, it will be available to everyone in the group, including those who don't contribute to its production.
3. There is no way (at least no practical way) to prevent noncontributing members from enjoying the benefits of the good.
4. The individual's contribution is a cost (e.g., in time or money or energy or emotional involvement) to that individual.

National defense is a public good. Many people must do their share (either by serving in the armed forces or by paying taxes to fund the defense budget). But once a country has a defense system, everyone will benefit from it, whether they contributed their share or not. And in many cases, it may not be clear whether a person does their share (they may be able to cheat on their taxes, for example). In fact there are many things that approximate being public goods.

Energy conservation is also a public good. Many people have to do their part if a meaningful amount of water or oil or gas is to be conserved. If enough people cooperate, they all enjoy the benefits of more and cheaper energy sources. This will be so even if a few people don't do their share, however, and in many cases it isn't possible to determine who does pitch in and who doesn't.

Other examples of public goods include population control, protecting the environment against pollution, preventing inflation, certain public health programs like immunization, and public education. Not everyone has to contribute for us to have public goods. We all expect there will be some slackers, but problems arise if there are too many. Public goods tend to generate two sorts of difficulties: the problem of free riders and the problem of assurance.

### **The Problem of Free Riders**

**Free riders:** people who benefit from a public good without helping to produce or maintain it

In some cities you pay for public transportation on the honor system. As long as everyone pays the system will work. But since so many people are involved, it will continue to work if a few people take a free ride at the expense of those who do not. This provides a nice metaphor for those who benefit from something without doing their share to produce it; these free loaders are said to be "free riders" on the efforts of those who do.

Rational individuals may easily conclude that they should take a free ride on the contributions of others because of the following (valid) bit of reasoning:

1. Either enough others will contribute for the good to be produced or else not enough will contribute.

2. If enough others do contribute, the good will be produced (my little input won't be missed), and so it is in my interest not to contribute.
3. If not enough people contribute, then the good will not be produced (despite the trouble I go to in doing my share), so it is in my interest not to contribute.

Since each person can reason in this way, many are likely to conclude that it is not in their interest to contribute. If so, the good will not be produced (even though it would be to the overall benefit of people to have it). The pattern of reasoning is the same as that in prisoner's dilemma; cooperation is necessary for the best results, but those who do not want to pay the price can enjoy the benefits of public goods by exploiting the cooperativeness of those who produce them (see figure 26.3).

		<b>Others</b>	
		Cooperate	Ride Free
<b>Me</b>	Cooperate	Enjoy Benefits, Costs of Cooperation	No Benefits, Costs of Cooperation
	Ride Free	Enjoy Benefits, No Personal Cost	No Benefits, No Personal Cost

Figure 26.3: Taking a Free Ride

### 26.2.2 The Assurance Problem

A person may be quite willing to forego a free ride and contribute *if* she could be assured that enough others would also contribute for the good to be produced. But since she realizes that many people may fail to do their share (perhaps because

they want a free ride), she lacks assurance about their actions. This provides an incentive for her not to contribute either.

The assurance problem becomes especially acute if some people are detected cheating. As distrust sets in, assurance plummets, and more people defect. This can trigger a downward spiral, since it will become more obvious that more people are shirking their responsibilities, and the motivations for others to do their share decreases.

### **The Tragedy of the Commons**

We will now consider four social dilemmas, some of which are all too real.

**The Commons goes down the Drain** In many nineteenth-century English villages, the families shared a modest parcel of land, the Commons, on which they grazed their cattle. They sold milk and butter and occasionally beef, and everyone made a reasonable living. But the more cattle a family had, the more it prospered, so each family had a motivation to add more.

As long as only one or two families added a couple of cows it didn't really damage the Commons. But every family could reason in this way, and they all added more cattle. At some point the Commons reached its carrying capacity; it was overgrazed and much of the grass died. As a result, cattle barely survived, some died, and profits dropped. As long as each family showed restraint, they all did well. When they each pursue their own interest, the result was worse for all.

**Voting** Most of us put a very high value on democracy; the ability to vote in free and open elections is one of the most important things about our country. But democracy only works if a reasonable percentage of eligible voters vote. Turnout is already low, typically less than fifty percent, but if it dropped to ten percent, or three percent, there would be little reason to view elected officials as legitimate representatives of the people.

So most of us want enough people to vote to maintain our democratic institutions. But it takes time and energy to go to the polls, and in a national election one person's vote has virtually no chance of making a difference. Of course the recent U.S. Presidential election, in which things went down to the wire in Florida, show that very small numbers of voters can have a very large impact.

Still, it is probably almost as likely that I will be run over by a steam roller on the way to the polls as it is that I will cast the winning ballot in a presidential election. And I am certainly more likely to be killed in a traffic accident than to determine the winner. So why *should* I vote? Why not sit back with a beer and take a free ride on the backs of those who take the trouble to go to the polls?

The deciding vote vs. the steamroller

**Campaign Costs** Raising the enormous sums of money needed to run for national office requires a great deal of time, and many people find it distasteful. Most candidates would prefer that everyone spent less on political campaigns. But each candidate also realizes that others are likely to spend more than they agree to spend, so there is pressure on everyone to spend more. This escalates, and we end up with our current situation, where big money plays a major role in who gets elected.

**Salary Caps** Most owners of professional football teams would prefer a situation in which no team paid vast, multi-million dollar salaries to their players. But efforts to cooperate often collapse, because sooner or later some team will spend a lot on some player who is just too good to pass up. Because the owners realize that they can't trust themselves, they placed a salary cap on what teams could spend.

### The State

Our largest cooperative endeavor is working together to maintain a society and government. Once a state is established there are various mechanisms that help keep it going. But in light of our recent examples we may wonder how a state could be established in the first place, and it is even possible to wonder how it keeps working, once it has been set up.

Of course some states do break down, but many do not. How is this possible? This was the first collective-action problem to occupy philosophers, and figures like Hobbes, Locke, and Rousseau used it as a starting point for their important discussions of the justification of the state (Locke's work, for example, was a major influence on the founders of this country).

In a large society like ours certain amount of free riding is bound to go undetected. The only way to prevent it would be to set up excessive monitoring and other oppressive measures, and this cure would be worse than the disease. This means that everyone will have numerous opportunities to take a free ride on the efforts of others, and so each person can reason as follows. Either enough people will cooperate to keep the state running (by being good citizens, obeying the laws, paying taxes, etc.), or else not enough will. If enough people do their part, then my little bit won't be missed, so I should take a free ride in any situation where I can get by with it. And if not enough people do their share, my tiny efforts won't keep the state running. Either way, I'm better off cutting corners whenever I can.

### Closer to Home

Prisoner's Dilemmas also arise in smaller groups. Consider a commune or large family that share a big house. All of them would prefer a clean house to a messy

one, and if each of them cleaned up the mess they made, the house would be spotless and neat. But the house is large, and if only one person cleans up their mess the house will remain almost as untidy as if they did nothing.

Of course they could clean up other people's messes too, but then others are free riding and this is likely to lead to resentment and an eventual refusal to do other's work for them. So there is some pressure for each person to leave the mess they make, with the result that the situation is worse for everyone.

Prisoner's dilemmas can occur any time trust is required. If two groups (e.g., protestants and catholics in Northern Ireland) are extremely suspicious of one another, they will be motivated to violate cease fires and other agreements for much the same reasons that our prisoner's were motivated to squeal. Trust, of course, is an issue in small groups, including groups of two, so we have worked our way back to the two-person case where we began.

### 26.3 How is Cooperation Possible?

When we focus on the reasoning in prisoner's dilemmas, it can begin to seem that social cooperation is impossible. Under some conditions it *is* very difficult to secure or maintain cooperation, but it clearly is possible in many situations because cooperative endeavors often do work very well. People do help one another, they do donate money to charities, they do pitch in, in all sorts of ways. Given what we have seen about social dilemmas, how is this possible? If we could understand why it works in some cases, it might be easier to find solutions in cases where it doesn't work so well.

In this section we will examine the several mechanisms that promote cooperation, but first a warning is in order. When we focus on energy conservation or pollution control cooperation looks like a good thing, and in such cases it is. But cooperation isn't intrinsically good; it all depends on the goals people are working to promote. A great many Germans worked together very efficiently to murder millions of Jews. As always, we have to think critically about each case.

Social dilemmas are a very mixed bag, and we shouldn't expect any single strategy to solve all of them. In some of our examples defection may be relatively easy to detect; in others it might be impossible. In some cases people know each other; in others they will never see each other again. In some cases coercion and sanctions are possible; in others they aren't. Some measures are relatively easy to implement; others would take a profound restructuring of society. But there are several devices that can promote cooperation in various conditions; they are not mutually exclusive, though some of them only make sense in certain settings.

1. **Coercion** We can force people to cooperate by imposing sanctions on those

Cooperation is not always good



who do not. The state, with its network of laws, police, judicial system, and prisons is well-suited to this task (some philosophers have argued that this, along with national defense, is the only real justification of the state).

2. **Rewarding By-Products of Cooperation** In many cases working together with others is rewarding in and of itself.
3. **Prudence and Prospect of Future Interaction** In some cases we will have to interact with people over a period of time. If they have discovered our propensity to take free rides, our reputation will suffer and they will be less likely to share benefits in the future. There may be a short-term gain from cheating, but it is offset by the long-term costs.
4. **Loyalty** If the members of a group know each other and care about one another, cooperation will often be easier to achieve.
5. **Moral Principles and Individual Ideals** Moral principles and individual ideals, like the principle not to lie or the ideal of being the sort of person who does their share, often prevent people from taking a free ride.

### 26.3.1 Coercion

Suppose that our prisoners *A* and *B* were members of the mob and knew that a contract would go out on them if they squealed. This changes the incentives or payoffs by raising the cost of defection to a point that the prisoners are likely to cooperate.

Many social dilemmas are solved by the coercive power of the state. Our tendency to obey authority, examined earlier in this module, is sometimes enough to get people to do their share. But self-interest is an extremely powerful force, and when it is involved authority alone isn't enough. Governments can require cooperation because they can enforce sanctions against those who are caught not cooperating. A major reason why people pay taxes is that they will be penalized if they don't; in effect the state forces us to pay for public goods by taxing us. The goal is to raise the cost of defection so that people will prefer to cooperate. There are many other examples of this sort of thing. One reason why companies don't dump pollution into the river, for instance, is that they will be fined if they are caught doing so.

Coercion can also be effective in some smaller groups. The workers in a factory may prefer the benefits of unionization (higher salaries, more job security) to their present situation. But they may also want to avoid the costs of being in the union themselves; one cost is union dues, but a higher cost is that the union might go

out on strike and strikers may lose their jobs. In states that allow “closed shops,” ones in which all the workers can be required to belong to the union, workers are coerced into joining and so they can’t take a free ride of the backs of those willing to strike.

The threat of coercion is one way of achieving cooperation, but there are many settings in which it won’t work. Coercion is possible at the international level (one country can place embargoes on the goods of another country, bomb it, invade it, and so on), but it is rarely effective. The absence of a threat of coercion is of course one reason why it is often difficult to secure cooperation on the international level.

In smaller groups coercion is often more effective, but it can’t be the whole story behind cooperation in such settings. It doesn’t explain why people often forego a free ride in those cases where they won’t be caught (e.g., some kinds of voluntary rationing). Nor does it explain why we often help others in situations where we wouldn’t be punished or criticized if we hadn’t lent a hand (e.g., giving directions to a stranger; sending money to poor children in another country). Finally, in some groups, for example a married couple, coercion would achieve cooperation only at the cost of undermining the group.

But even when coercion is possible, it is not an attractive way to promote cooperation. Big, coercive government is very dangerous. Furthermore, if people only cooperate because they have to, they are much more likely to seek opportunities to cheat. It would be better all the way around if people cooperated because they wanted to.

### **26.3.2 Positive By-Products of Cooperation**

In some cases cooperative endeavors are rewarding in and of themselves. Many people take pleasure in working alongside others for something they view as a good cause. Striving to promote a common end, like stacking sand bags by a river that’s about to flood the town or banding together to fight for social justice, can generate friendships and feelings of solidarity. Even intrinsically unpleasant acts like giving blood can seem rewarding when we do it to help others who need blood.

In many cases social by-products do a lot to promote cooperation, but they can’t explain all cases. Even people who think we need taxes experience little solidarity with fellow taxpayers as they write out their checks on April 15. Few of us feel warm fellowship with others when we turn our thermostats down to conserve gas or toss away our aerosol sprays because they might deplete the ozone layer. Moreover, while the positive side-effects of cooperation often do help maintain cooperative enterprises, they don’t explain how those enterprises get off the ground in the first place.

### 26.3.3 Prudence and the Prospect of Future Interaction

The members of small groups are likely to know each other and to interact over a period of time. Here pressures of conformity become stronger, and the worry about one's reputation increases. If people come to think of me as a social parasite, my reputation will suffer. Many people find this aversive in itself. But it also means that others will be less willing to trust me or help me in the future. So often it is better to forego a short-run benefit (getting a free ride) in favor of the long-term benefits of cooperation. Free-riding may be in my *short-term* self-interest, but it is often in my *longer-term* self-interest to do my share.

We often do things we don't particularly like doing now in order to enjoy benefits later. But we are human, and sometimes we go for good old instantaneous gratification. Something *right now*, a cigarette, another piece of pie, a free-ride on the efforts of others, often wins out over our longer-run best interests. So foresight and prudence have their limits in the promotion of cooperation.

Worries about the prospect of future interaction diminish as the size of the group increases. In the small rural community where I grew up, everybody knew everybody else, but large cities are much more anonymous. In large groups it is easier to take free rides without detection, and even if others notice, the chances that we'll ever meet again are slim. There is a good deal of evidence that smaller groups cooperate better than huge groups, and this is one reason why.

#### Iterated Prisoner's Dilemmas

People who will interact with each other over time are not in the same position as our prisoners *A* and *B*. With frequent contact, we learn which people can be trusted, and we adjust our behavior accordingly. Some aspects of this setting can be modeled by **iterated prisoner's dilemmas**.

Here 'iterated' means 'repeated', so that two (or more parties) play a series of prisoner's dilemma games. After a few rounds, each party has some idea about how trustworthy the other parties are. Groups of people can do this, and social scientists have studied their behavior when doing so. You can also do a computer simulation to see which strategies work best.

It turns out that the strategy that will win under many conditions is *tit-for-tat*. If the other person cooperates on one trial you cooperate on the next (thus rewarding cooperation). And if the other person defects on one trial you defect on the next (thus punishing defection). Why do you think this strategy works?

The fact that *tit-for-tat* wins out under certain conditions may help explain why people would cooperate with others when they knew they would be interacting with them in the future. But we often help strangers that we know we'll never see again,

so it can't be the whole story.

### 26.3.4 Loyalty

Suppose that our prisoners *A* and *B* knew each other well and cared about one another's welfare. Perhaps they are brothers or old friends. Each would then have reasons to trust the other and to care about his welfare, and this would make them more likely to cooperate. More generally, groups that have common aims and shared understandings can often foster feelings of trust and solidarity. As groups become larger and more diverse, this becomes more difficult.

### 26.3.5 Moral Principles and Individual Ideals

#### What if Everybody did that?

Ann: I know it's election day, but I'm sort of tired. I'm just not going to vote.

Betty: But what if everybody did that?

Ann: But everybody doesn't.

Betty: What if everybody reasoned that way? Then *nobody* would ever do the socially responsible thing.

But Betty's responses simply invite the now-familiar litany. Either (almost) everyone will do the socially responsible thing or they won't. If they do, then my contribution won't be needed; if they don't, my efforts to be a good citizen will be in vain. As long as we think in terms of short-term payoffs for ourselves, the question "What if everybody did that?" isn't going to faze us. We need reasons, other than narrow self-interest, to be moved by Betty's line of argument.

#### Principles and Ideals

In fact most people *are* motivated by things other than narrow self-interest. In the process of socialization members of a culture internalize moral principles and individual ideals, and these exert an extremely powerful influence on behavior. For example, codes of honor played a central role in many cultures, particularly strongly hierarchical ones, and many people preferred death to dishonor. Indeed, throughout history people have been willing to die for a cause.

In our original scenario the prisoners were chiefly motivated by a desire to spend as little time in prison as possible. But suppose instead that *A* and *B* were political prisoners struggling for a cause they thought extremely valuable. In this

case one or both might be willing to make large sacrifices, even give up their life, to remain loyal to their cause and their comrade.

Altruism also leads many people to help others even when it is not in their narrow self-interest to do so. When the Nazis occupied France during World War II, many French people heroically risked their lives to hide Jews that they hadn't even known. We do have a capacity to act in accordance with principles and from feelings of altruism, as well as out of narrow self-interest, and these often promote cooperation.

Most of us, for example, believe that lying is wrong. This principle will make us more likely to cooperate in cases when we have promised to do so. Our self-image may also be that of a person who does their share, who keeps their bargains, whether other people do so or not. Even if you could cheat without being detected, you wouldn't do so because you think that it's wrong, or because you don't want to be the *sort* of person who cuts corners or takes the easy way out.

Suppose that my chief motivation for cooperating is mere prudence; I don't want to be *caught* cheating because people wouldn't treat me very well in the future. Since my motivations are to avoid suffering the consequence of detection, it makes as much sense for me to seek subtle ways of cheating as it does for me to cooperate. But if I am motivated by the belief that cheating is wrong, I won't do this. It would be a mistake to underestimate the power of norms and ideals, but as history makes lamentably clear, they are too weak to sustain all cooperative endeavors (including living together in peace) without help from other sources.

## 26.4 Fundamental Values: Clashes and Tradeoffs

### Basic Values

Members of different cultures often have at least somewhat different values. We will consider this below, but let's begin thinking about ourselves. In a large, heterogeneous country like the United States, there are many disagreements about values, but there are some basic values that most of us share. Many of the most basic disagreements in our society involve clashes among these values. Often these values reinforce each other, but there are also tensions among them. In some cases—the hard ones—you can have more of one value only by having less of another. Consider the following values:

1. Freedom or liberty
2. Majoritarianism (majority rule, i.e., democracy)
3. Equal opportunity
4. The need for security

## 5. Community moral standards

### Basic Tradeoffs

Virtually everyone in our country believes that democracy, the rule of the many, is a good thing. But democracy can be in tension with other values, most obviously individual rights and liberties. A majority can tyrannize a minority just as much as a dictator can. For example, until the 1960s poll taxes and other public policies made it almost impossible for African-Americans in many parts of this country to vote.

A chief purpose of a constitution, a Bill of Rights, and judicial review (review of lower court decisions by higher courts) is to guard against a tyranny of the majority. People have certain fundamental rights that cannot be violated, even if the majority wants to. But where should we allow majority rule, and where do individual rights override it?

Freedoms or liberties or rights are in tension with other values. For example, many of us might be more secure if the government relaxed the rules on search and seizure or wire tapping. Indeed, there are even tensions among liberties themselves. For example, freedom of expression (freedom of speech and freedom of the press) is absolutely fundamental. Without it our other freedoms could be abused without most of us even hearing about it, and that would make such abuses easier and more frequent.

Still, most of us think that there should be some limits on freedom of the press. In this century, as technologies like wire tapping and video taping have advanced, we have become increasingly sensitive to people's rights to privacy. For example, reports in the media wrecked Richard Jewell's life, and many people think that the intrusion into Princess Di's life played a role in her death. Many people also believe that names of rape victims should not be published. The difficult question is how to balance freedom of expression with rights to privacy.

Most of us think there should also be other limits to freedom of expression. We favor laws against libel and slander, child pornography, and publishing manuals explaining how to construct nuclear weapons. Freedom of the press can also conflict with a right to a fair trial. If we are constantly hearing about a terrible crime, it can make it much more difficult for the defendant to receive a fair trial. But as debates over censorship of the internet show, balancing such considerations is never easy.

Most of us also think that everyone should have the opportunity at a decent life. But children who grow up in the ghetto and walk past drug dealers and metal detectors to get to class clearly don't have the opportunities many of the rest of us do. They deserve help, but the only way to provide it is by the redistribution of money (in the form of taxes to fund welfare programs, public health facilities,

school lunches, and the like). Again there is intense debate about how to balance values.

Finally community values can conflict with other fundamental values. Laws against prostitution and selling drugs place restrictions on the rights of adults to engage in free business transactions, and censorship laws restrict freedom of expression.

### Discussion Topic

Tom Palance is trapped in a cave with a very hungry and sick child that will die if she doesn't eat soon. They both know that they will be rescued in a matter of hours, but the child will probably starve to death before that happens. Tom has a number of egg salad sandwiches that will spoil if they aren't eaten in the next couple of hours. He eats a sandwich and a half then, full, tosses the rest into a deep lake rather than giving one to the hungry child.<sup>1</sup>

1. Is his action wrong?
2. Why or why not (what *principles* might lie behind our judgments about whether or not it is wrong or not)?

Now suppose that someone proposes the following argument:

Most of us are actually in a situation that is very similar to Tom's. We could, at minimal cost to ourselves, do things that would greatly improve the lot of starving people in various places. Indeed, even just the things we waste, or the money we spend on things we don't need, could do a lot to help people who are really suffering through not fault of their own. So we are really no better than Tom.

3. Your task is to evaluate this argument.
  - (a) The argument turns on an analogy. Just what *is* the analogy (what two things are said to be analogous)?
  - (b) What relevant *similarities* are there in the two things or cases?
  - (c) What *irrelevant similarities* are there in the two cases?
  - (d) What relevant *differences* are there in the two cases?
  - (e) What *irrelevant differences* are there in the two cases?
  - (f) What principle (or principles) do you think lie behind this argument?
  - (g) *Is this a good analogy?*

---

<sup>1</sup>A good discussion of public goods and free riders can be found in Mancur Olson's classic *The Logic of Collective Action: Public Goods and the Theory of Groups*, Cambridge MA.: Harvard University Press, 1965. For an excellent discussion of many of the themes in this chapter see Russell Hardin's *Collective Action*, Johns Hopkins University Press, 1982. Further references to be supplied.

- (h) Could the argument be rephrased without using the analogy? What would be lost if it was?
- (i) Could the story about Tom be revised in any ways that would make the analogy better?
- (j) Should we have a *law* that requires people in Tom's situation to help the child?

## 26.5 Chapter Exercises

1. *Group Exercise:* Suppose you are working on a group exercise in class and that your personal grade for the project will be identical with the grade earned by your group. But part of the work for the project requires that each person spend some time gathering information in the library before the next meeting when the group will complete its project.

Is there an assurance problem in this setting? Is there a danger of free riders in this setting? To what extent do the features of this situation mirror the problems involving public goods (go through the conditions for a public good one by one)? To what extent are they different? What measures might minimize the problem of free riders?

2. Explain why each of the following is a public good. How might people try to take a free ride, enjoying its benefits without contributing, and how might this be prevented?

- 1. population control
- 2. protecting the environment against pollution
- 3. preventing inflation
- 4. various public health programs like immunization
- 5. public education.

3. In the chapter on groups we examined social loafing (p. 486). How is it related to free-riding? In what ways are they different? In what ways are they the same? What ways of fostering group cohesiveness decrease free-riding?

4. When I moved to Norman there was a shortage of apartments and prices were high (at least they seemed high to me). Many builders noted the shortage and began construction of new apartment complexes. In a few years there were too many apartments; landlords had to maintain vacant apartments, and many actually offered to lower rents to entice people to stay. They would have preferred a situation where the apartments were all filled and rents were higher. Why didn't this happen? Is this a Prisoner's Dilemma? If so, why; if not, why not?



5. Owners of professional football teams might impose salary caps because they cannot trust themselves to stick to an agreement to pay lower salaries. Can you think of other situations like this? We sometimes vote to have taxes to support government activities or dues to support an organization. In what ways is this similar to the owners' use of salary caps.
6. You live in a large house with a group of other people. All of you would prefer a clean house to a messy one, and if each of them cleaned up the mess they made, the house would be clean. But the house is large, and if only one person cleans up their mess the house will remain almost as messy as if they did nothing. As a result no one ever seems to clean up anything. How might you use each of the four measures discussed above (p. 530; coercion, setting things up to yield rewarding social by-products, prudence, and moral principles and ideals) to do this. Which of the four do you think would be most effective? Which least?
7. You have been hired to promote more energy conservation in the state of Oklahoma. How might you use each of the four measures discussed above (coercion, setting things up to yield rewarding social by-products, prudence, and moral principles and ideals) to do this. Which of the four do you think would be most effective? Which least?
8. You have been hired as a consultant to a project promoting population control in an over-populated third-world country. Your thoughts here will have to be somewhat tentative, since you don't know the culture, but your co-workers do. How might you use each of the four measures discussed above to do this? Which of the four do you think would be most effective? Which least?
9. Tensions have been mounting between two nuclear super powers. They have become increasingly distrustful of each other, and each now fears a surprise attack. In high level meetings in both countries, the defense ministers argue that the only sensible option is a quick, surprise strike of their own. What reasoning might those in the meeting go through at this point? What might be done to avert the worst outcomes?
10. Over-fishing of various fishing areas has dramatically depleted the number of fish. In what ways does this situation resemble the Tragedy of the Commons? Can you think of any other contemporary counterparts of the Tragedy of the Commons?
11. Why *should* we vote?
12. We began with a simple game that has implications for many social situations. We'll conclude with another game called *chicken*. It is named after the game where

two people (mostly teenage males) drive their cars straight at each other and the first one to swerve out of the way is a chicken; there is a gripping depiction of this game in the movie *Rebel without a Cause*. In Figure 26.4, each person prefers  $a$  to  $b$  to  $c$  to  $d$ . Can you say what these outcomes amount to?

		<b>B</b>	
		Swerve	Not Swerve
<b>A</b>	Swerve	$b, b$	$c, a$
	Not Swerve	$a, c$	$d, d$

Figure 26.4: Playing Chicken

1. Describe what  $a$ ,  $b$ ,  $c$  and  $d$  stand for in a standard game of chicken.
2. The structure of this game differs from the Prisoner's Dilemma. Here we hope to get the other person to cooperate (swerve first) by trying to convince them that we will *not* cooperate. What real-life situations exhibit the structure of this game?
3. What dilemmas and traps can such situations generate?

## **Part X**

# **Representation and Recognition**



## **Part X. Representation and Recognition**

In chapter 27 you will learn to draw diagrams that make it easier to reason about the logical structure of many arguments and the probabilistic structure of many problems. Pictorial representations often allow us to see—literally—what the solution to a problem is.

In chapter 28 we turn to a few very general tools for thinking more critically and accurately and discuss the best ways of learning to reason more effectively. Among other things we discuss ways to become sensitive to cues that a particular cognitive tool is relevant in situations outside the classroom.



## Chapter 27

# Diagrammatic Reasoning: Using Pictures to Think

**Overview:** In this chapter you will learn to draw pictures and diagrams that make it easier to reason about certain sorts of problems. We begin with diagrams of the logical structure of arguments, then move on to diagrams that are useful in various types of inductive reasoning. Such diagrams often allow us to see—literally—the key aspects of a problem, and sometimes even its solution, in a flash. We focus on diagrams that can be drawn easily, often in just a few seconds on the back of a napkin.

### Contents

---

<a href="#">27.1 Using Vision to Think</a>	545
<a href="#">27.2 Picturing Logical Structure</a>	547
<a href="#">27.3 Picturing Probabilistic and Statistical Structure</a>	553
<a href="#">27.4 Chapter Exercises</a>	570
<a href="#">27.5 Appendix: Inverse Probabilities and Bayes' Rule</a>	571

---

## 27.1 Using Vision to Think


### Pictures can Make Problems Easier

There is an ancient saying that a picture is worth a thousand words. Is it true? Well, it all depends on the particular picture, the words, and the problem we want

Pictures make many problems easier

to solve. But many problems that seem hopelessly confusing when described in words often become surprising clear when we find a way to represent them in a diagram. Indeed, with the right picture a problem may pretty much solve itself; the answer just jumps out at us.

We have all heard about first-cousins twice-removed, second-cousins one-removed, great uncles, and the like, but most of us aren't very clear about what such relationships amount to; if we draw a family tree, however, things begin to fall into place. And the detailed workings of supply and demand and equilibrium in Econ 101 leave our heads spinning, but they begin to make sense when we learn to draw supply and demand graphs or curves. Indeed, there are many graphical or pictorial representations that make thinking easier: blueprints, seating charts, maps, flow diagrams and, increasingly, computer graphics and animations.



Remember Linda, the outspoken, bright, philosophy major (p. 312). We asked whether it is more probable (1) that she is a bank teller, or (2) that she is a bank teller who is active in the feminist movement. If we give the second answer we commit the conjunction fallacy, but the problem can be confusing, and it may be difficult to see why this answer involves bad reasoning.

Once we draw a picture, though, the answer pops out at us. Here the crosshatched area where the circles overlap represents the set of bank tellers who are also feminists. Clearly this area cannot be larger than the entire circle on the left (which represents bank tellers in general) *or* than the entire circle on the right (which represents feminists in general).

### Why are Pictures so Useful?

A large portion of the human brain (the visual cortex) is devoted to vision, and we are highly visual creatures. So representing problems pictorially plays to our strengths.

One reason it is useful to present information pictorially is that this format helps compensate for limitations on working memory. We can't focus directly on very much information at a time; our attention and working memory are very limited. You can keep a six or seven digit number (like a phone number) in your working memory as long as you keep repeating it to yourself, but once you get to nine or ten digits it's very difficult and with twenty or so it's hopeless.

Reasoning frequently requires us to focus on a good deal of information at once, and often we just can't keep everything straight. Pictures can help. For one thing, they allow us to keep information that needs to be used together close together, so that we don't have to keep searching around for the information we



need.

Diagrams allow us to represent abstract relationships between individuals or sets with simple geometrical relationships like the inclusion of one circle in another (as in Figure 27.1 on page 549) or the overlap of a pair of circles (as in the above diagram of feminist bank tellers). We can also represent facts about percentages and proportions (which is often a good way to think about probabilities) by the relative size of different parts of a diagram, as we do in a bar graph or pie chart (see Figure 27.8 on page 556).

Such representations are useful because we human beings are very good at recognizing geometrical relationships like overlap, inclusion, and relative size. When we exploit such structure in a diagram, the information is organized in a way that we can take in at a glance, and often we can draw visual inferences almost automatically.

Diagrams are also useful for exploration and communication. When we draw a diagram we can tinker with it, erasing the circle here, adding a dot over there, patching it up by trial and error as we grope our way to clarity. And in many cases it is easier to communicate an idea with a diagram than with words.

Different situations and problems call for different sorts of diagrams, and some problems aren't usefully represented by diagrams at all. Furthermore, diagrams, like all representations, can distort, muddy, confuse. Still, when we *can* draw a good diagram it often cuts a very confusing problem down to a manageable size.

In this chapter we will get a feel for the sorts of problems that can be tackled with diagrams and learn something about which sorts of diagrams are useful where. The goal is to add diagramming to the set of cognitive tools that you can actually *use*, so we will focus on simple diagrams, ones you can often construct in a matter of seconds on the back of a paper napkin.

Don't make pictures any more complicated than necessary

## 27.2 Picturing Logical Structure

Just as a good blueprint can help us visualize the *spatial layout* or *structure* of a house, a good diagram can help us grasp the *logical layout* or *structure* of a sentence or an argument. Indeed (as we will see below), in some cases when we draw the picture of the premises of an argument, the conclusion is automatically drawn in the process.

Consider the two very simple arguments below:

The two arguments have very different content or subject matter (human mortality vs. numbers), but they also have something very important in common. Both are valid, and they are valid *for the same reason*. As always, validity is based on form

P1. All humans are mortal. <u>P2. Wilbur is human.</u> So Wilbur is mortal.	P1. All even numbers are divisible by 2. <u>P2. 18 is an even number.</u> So 18 is divisible by 2.
---	--

or structure, rather than content or subject matter, and these two arguments have the same logical structure.

We can represent this common structure by erasing the content and just leaving the logical skeleton behind:

P1. All _____ are _ _ _ _ _ .	
<u>P2. <math>x</math> is a _____ .</u>	
So $x$ is a _ _ _ _ _ .	

No matter how we plug content into these blanks, as long as we plug the same word or phrase into for the same placeholder each time it occurs and the result is a grammatical sentence, the resulting argument will be valid.

In the best case you can draw a picture of the premises and you will find that your picture, without any further additions, also includes a picture of the conclusion. This is the case in some of the simpler examples of deductively valid arguments. Indeed, in these cases teachers often draw a picture to convince their students that a given argument format is valid.

It will help if we imagine constructing a picture of this argument structure step by step (see Figure 27.1 on the next page for the full picture). We begin by drawing the first premise, putting a smaller circle representing humans entirely inside a larger circle representing creatures that are mortal (in the subfigure at the left).

Next we draw the second premise by placing a letter representing Wilbur in the second circle (in the middle subfigure).

Finally, we need to draw the conclusion. But we find that in drawing the two premises we already *drew* the conclusion. This occurred because the argument is valid, and as we saw earlier, the information in the conclusion is already contained in the premises. Since the information in the conclusion is already in the premises, our representation of the premises will automatically contain a representation of the conclusion.

The left and center diagrams in Figure 27.1 on the facing page are already fairly abstract, but in the diagram on the far right we remove the very last bit of content. This leaves just a picture of the argument's form. It is now fairly easy to see that—and why—any argument with this form must be valid.

Some sentences and arguments do not lend themselves to a graphical representation, but many do and for those that do, pictures can be very helpful. Of course this is an easy example, and we don't really need a digram to understand it. But




Figure 27.1: Picture of the Structure of an Argument

pictures really come into their own when we turn to more difficult or complex problems that can be hard to grasp without pictures (we will see an example below with the card-selection problem)

### Conditionals and Conditional Arguments

#### Necessary and Sufficient Conditions

Thinking in terms of frequencies or proportions not only helps us get clearer about probabilities. It can also sometimes help us reason about conditionals. Take the conditional

If Fido is a dog, then he's a mammal.

This sentence is about a particular dog, Fido, rather than about groups of any sort. But many (not all) conditionals like this are true (or false) because of facts about groups, here the fact that *all dogs are mammals*.




Figure 27.2: Picture of Necessary and Sufficient Conditions

We can represent *this* statement in with a digram like that in Figure 27.2 on the page before, and then we can use it to help us in our reasoning. We can place the dot representing Fido in the *dog* circle, and it's immediately clear that he's in the *mammal* circle.

So far things are pretty obvious our diagram also shows us why various, more complicated sentences are true or false. For example, if we place Fido's dot anywhere outside the mammal circle, say out at location *Z*, there is no way he could be in the dog circle. So the conditional

If Fido is not a mammal, then he's not a dog.

is also true.

Now let's return to Figure 27.2 to see how necessary and sufficient conditions work.

- If something is *inside* the smaller, dog, circle, then it must be inside the larger, mammal, circle. Being in the smaller circle is *sufficient*, all that it takes, for being inside the outer circle. Since *X* is a dog, *X* is a mammal.
- If something is *outside* the outer circle there is no way it could be inside the inner circle. As *Y* shows, being in the larger circle isn't sufficient for being in the smaller one, but it is necessary.

In those cases where we can think of conditionals in terms of sets of things, the set (e.g., dogs) mentioned in the antecedent is the inner circle and the set (e.g., mammals) mentioned in the consequent is the outer circle. This makes it clear pictorially why antecedents are sufficient conditions and consequents are necessary conditions.

This is possible in many instances where it may be less obvious. Consider:

If my alarm breaks tomorrow, I'll miss my flight.

We can think of this as saying: on all occasions in which my alarm didn't go off I would miss my flight. This is so, anyway, on all the normal occasions that are likely to arise. We can then draw a picture of the familiar sort, with the circle containing the occasions where my alarm breaks inside the circle of the occasions on which I miss my flight.

When we commit the fallacy of *affirming the consequent* (p. 54), we are arguing as though being in the larger circle were enough to ensure being in the smaller one. But *Y*'s location in Figure 27.2 on the preceding page shows this is wrong. And when we commit the fallacy of *denying the antecedent* (p. 55), we are arguing as though not being in the smaller circle were enough to ensure not being in the larger one. Again *Y* shows that this is invalid. It's all pretty clear when we stick to dogs and mammals, but we'll now look at a case that many people find confusing.

### The Card Selection Task

One of the exercises in the chapter on conditionals went like this:

Wilbur has a pack of cards, each of which has a letter [either a consonant or a vowel] on one side and a number [either even or odd] on the other. Some of the cards lying flat on a table (Figure 27.3). Which cards should you turn over in order to test the hypothesis:

Hypothesis: If a card has a vowel on one side, then it has an odd number on the other.

Hint: only two of the four kinds of labels are relevant. The problem can be con-




Figure 27.3: Cards in the Selection Task

fusing, but once we draw a picture that represents the hypothesis, things become easier. The hypothesis means that if *any* card has a vowel on one side then it has an odd number on the other, so it's claiming that *all* cards with vowels on one side have odd numbers on the other. And we know how to draw a picture of this, which we do in Figure 27.4.




Figure 27.4: Testing the Hypotheses in the Selection Task

Note that the hypothesis only says that all vowels have odd numbers on the other side. It does *not* say anything about what odd numbers have on the other side, and it does *not* say anything about what consonants have on the other side. This means that there is no point in flipping either of these sorts of cards. Suppose that I do flip a card with an odd number. Will it matter what is on the other side?

No, because our hypothesis doesn't make any prediction about that. No matter *what* is on the other side, it's compatible with the hypothesis.

But what about cards that have even numbers on them? Those are the ones lying outside the big circle in Figure 27.3, out in areas like those occupied by the three stars (\*\*\*). We can see that cards out here cannot be vowels. So we *do* need to flip cards out here (which all have even numbers face up). If a card out here has a vowel on it, the hypothesis we are testing is false.

We can also use the diagram to explain the answer in terms of necessary and sufficient conditions. The hypothesis says that having a *vowel* on one side is *sufficient* for having an odd number on the other. So we need to flip each vowel card, and if we find even one without an odd number on the other side, the claim is false.

Similarly, the hypothesis says that having an *odd* number on one side is a *necessary* condition for having a vowel on the other. So if we find just one even card with a vowel on the other side, the claim is false. We can see from our picture that the hypothesis requires that if a card is not odd (= even), then it is not a vowel. If we find a card, say \*\*\*, that is not odd but is a vowel, the hypothesis is false. The relevant argument pattern here is denying the consequent.

There is some evidence that people are better at problems like this when they involve rules and obligations. We do much better if you ask us which people a policeman must check to see if they everyone is obeying the rule:

Rule: If you buy beer you must be at least twenty one.

What is the answer? Why might we do better with problems like this?

### Picturing Validity




Figure 27.5: Picturing Validity

An argument is valid just in case it is impossible for all of its premises to be true while its conclusion is false. This means there is no way all the premises could

be true and the conclusion false. So in every possible situation in which all of the premises are true, the conclusion must be true as well.

We can draw a picture of this by representing the set of all the possible situations in which the premises are true by the inner, smaller circle in Figure 27.5 on the preceding page and the set of all the situations in which the conclusion is true by the larger, outside circle.

By contrast, if an argument is *invalid*, there is some possible scenario in which the premises are true and the conclusion is false. Of course the definition of validity allows for the conclusion to be true in additional situations in which one or more premises are false.

### 27.3 Picturing Probabilistic and Statistical Structure

We can use simple diagrams to illustrate various probabilistic concepts. Indeed, we can often use them to show why certain probabilistic rules are correct. Some of the diagrams appeared in earlier chapters, but it will be useful to collect them together here and to add a few new tricks.

#### Probabilities of Disjunctions

We begin with two diagrams we encountered earlier, to remind ourselves how pictures can help us think about probabilities. Recall that a disjunction is an *either/or* sentence. What is the probability that a disjunction,  $A$  or  $B$ , when the two disjuncts,  $A$  and  $B$  are *incompatible*?




Figure 27.6: Probabilities of Disjunctions

We can represent the situation with the left subfigure in Figure 27.6. To know how much area is covered by either  $A$  or by  $B$  we would simply add the two areas together. We represent the fact that the disjuncts are incompatible by a picture in

which the two circles representing them do *not* overlap, so it is easy to see that we didn't count any portion of the area twice when we added things up.

By contrast, if the disjuncts  $A$  and  $B$  are compatible, then they can occur together. We represent the fact that they are compatible by a picture in which the two circles representing them *do* overlap. And this makes it easy to see that if we simply add the area of  $A$  to the area of  $B$ , we will count the overlapping (crosshatched) area twice (once when we consider  $A$  and a second time when we consider  $B$ ). The solution is to subtract the probability that  $A$  and  $B$  both occur so that this area only gets counted once.

We now turn to more difficult problems that are much easier to solve with pictures than with numbers.

### Wilbur Fails his Lie Detector Test

The police have word that there is a large conspiracy to sell stolen software. In an effort to put it out of business they round up a hundred random suspects and give them lie detector tests. Suppose, hypothetically, that the following numbers are right.

- The machine says a person is lying in 90% of the cases where they really are lying.
- The machine says a person is lying in 20% of the cases where they are *not* lying.
- 10% of the suspicious characters in town really are guilty.

Wilbur failed his test; the machine says he lied. What is the probability that he is guilty?

It is tempting to reason like this: the test is 90% reliable, so there must be about a 90% chance that Wilbur is guilty. The 90% here is the conditional probability of someone's failing the test if they are guilty.

But this is *not* what we want to know. We want to know the *opposite* or *inverse* of this; we want the probability that someone is guilty given that they failed. The first figure is  $\Pr(\text{Fails} \mid \text{Guilt})$ , but we want to know the inverse of this,  $\Pr(\text{Guilt} \mid \text{Fails})$ . The two probabilities can be very different; compare:

$$\Pr(\text{Male} \mid \text{Pro football player}) \neq \Pr(\text{Pro football player} \mid \text{Male}).$$

In the appendix to this chapter we will see how to use our rules to calculate the probability of Wilbur's guilt. But a simple picture can be used to do solve the problem, and it is much easier to understand.



Not Guilty									Guilty
									-
									-
									-
									-
									-
									-
									-
									-
-	-	-	-	-	-	-	-	-	-
-	-	-	-	-	-	-	-	-	-

Minus (-) = failed test; empty cell = passed test

Figure 27.7: Probability of Guilt Given Test Results

Figure 27.7, which is drawn to scale, represents the relevant information. Remember that there are 100 people in our group.

- 90% of them are *not* guilty, and they are represented by the wide column on the left of the diagram. Ten percent *are* guilty, and they are represented by the much thinner column on the right.
- 90% of the guilty people fail the test, and we indicate this with a minus, -, in 90% of the rows in the thin column.
- 20% of the people who are not guilty fail the test (perhaps they are nervous), and we indicate this with a minus in 20% of the area of the wide, non-guilty column.

Now we just need to consult the figure and check the percentages. When we do, we find that 9 of the 10 guilty people fail the test but 18 of the innocent people also fail (18 is 20% of the 90 people who are innocent). So only one person in three who fails the test is guilty. If we have no additional information to go on, the probability that Wilbur is a conspirator is  $1/3$ . It is important that you count, literally, the number of minus signs in the figure to verify this.




Figure 27.8: Pie Chart of Lie Detector Test

Similar points apply in many other situations. For example, we may want to know about the probability of having breast cancer given a small lump and a certain result on a mammogram. Problems like these require us to *integrate* knowledge about the present case (e.g., what the lie detector says) with prior information about *base rates* (proportion of guilty people among the suspects). We often focus too much on the present information (Wilbur failed) and ignore information about the base rates of guilty people among the suspects (only 1 in 10).

There are often various ways to display information. We could also depict Wilbur's plight with a pie chart in which each of our four groups is represented by a slice of the pie (Figure 27.8; it is drawn to scale, so that 10% of the pie represents 10% of the group). Such diagrams are often easy to read, but it takes a bit more work to construct them than it does rectangular diagrams.

You should use whatever sorts of pictures you find most helpful, but it is important to construct quick and simple figures; otherwise they will take so much work that (if you are like most of us) you will rarely put in the effort to construct anything at all.

### Exercises

1. Draw a picture like that in Figure 27.7 on the page before and determine the probability of Wilbur's guilt if all the numbers remain the same except that the probability of failing the lie detector test if you are not guilty is only 10% (instead of 20%, as above).
2. Draw the picture and determine the probability of Wilbur's guilt if all the numbers remain the same except that the base rate of guilty people among the suspicious people is 20% (instead of 10%, as it was above).

3. The staff at the Suicide Prevention Center know that 2% of all people who phone their hot line actually attempt suicide. A psychologist has devised a quick and simple verbal test to help identify those callers who will actually attempt suicide. She found that

- a. 80% of the people who will attempt suicide have a positive score on this test.
- b. but only 5% of those who will not attempt suicide have a positive score on this test.

If you get a positive identification from a caller on this test, what is the probability that he would actually attempt suicide?

4. Suppose that we have a test for the HIV virus. The probability that a person who really has the virus will test positive is .90 while the probability that a person who does *not* have it will test positive is .20. Finally, suppose that the probability that a person in the general population has the virus (the base rate of the virus) is .01. Smith's test came back positive. What is the probability that he really is infected? Construct a diagram to determine the answer (we will work it out by the numbers in the appendix, so you can check your answer against the answer there).

### Correlations Revisited

The discussion of correlations in Chapter 15 emphasized pictures, so we will simply recall the basic points here. We saw that the easiest way to understand the basics of correlation is to use a diagram like Figure 27.9 on the following page. It depicts four hypothetical relationships between smoking and getting lung cancer.

The fact that the percentage line is *higher* in the smokers column than it is in the nonsmokers column in the top two diagrams in Figure 27.9 on the next page indicates that smoking is positive correlation between being a smoker and having a heart attack. It is the *relationship* between these two horizontal lines that signifies a positive correlation. And the other diagrams work in the same way.

It bears repeating that you do *not* need to know exact percentages in order to draw most of these diagrams. You only need to know which column has the higher percentage, i.e., the higher horizontal line.

### Probability Trees

Rectangular diagrams are very clear, but only work well when we are considering two factors or variables or outcomes. Figure 27.10 on page 559 shows how to represent three variables; here the outcomes of three successive flips of a fair coin, with each flip is a separate variable.




Figure 27.9: Diagrams of Correlations between Smoking and Lung Cancer




Figure 27.10: Tree for Three Variables

In the tree the numbers along each path represent the probabilities of the outcomes. The probability of a heads on the first flip (represented by the first node of the top path) is  $1/2$ , and the probability of a second head as we move to the right along that path is also  $1/2$ . There are eight complete paths through the tree, and to assess the probability of any of them (e.g.,  $H_1T_2H_3$ ) we simply multiply the numbers along the path. Since the coin is fair, the probability for each of the eight outcomes represented by the eight full paths is  $1/8$ .

Trees can also be used when we have fewer than three variables. In Figure 27.11 we represent the data about Wilbur's lie detector test that we displayed above in a rectangular diagram (p. 555) in a tree format. Here we represent the probabilities as labels on the branches.

For example, the branch beginning with Suspects and running to Guilty is labeled 10% to indicate that 10% of the suspects are guilty. And the bottom branch from Suspects to Not Guilty is labeled 90% to indicate that 90% of the people rounded up are innocent. The additional labels represent the additional probabilities. For example, the top right branch is labeled 10% to indicate that 10% of the guilty people manage to fool the machine and pass the test.

The arrows point to the numbers of people who fail the test. There are 9 guilty people who fail plus 18 innocent people who fail, for a total of 27 failed test. Of these 27 failures 18, or  $2/3$ , are innocent. So if we have no additional information,

our best guess is that the probability that Wilbur is guilty is  $1/3$ .





Figure 27.11: Tree Diagram of Wilbur's Lie Detector Test

### Exercises

1. Represent the information about the officials at the Suicide Prevention Center above (p. 557) as a tree and use it to solve the problem.
2. Represent the information about the HIV test above (p. 557) as a tree and use it to solve the problem.

### Brute Enumeration of Percentages

#### Pictorial Solution to Monty Hall Problem



In the appendix to this chapter we will solve the Monty Hall Problem using probability rules (p. 578), but most people find a picture of the situation more illuminating. Recall the problem: there are three doors in front of you. There is nothing worth having behind two of them, but there is \$1000,000 behind the third. Pick the right door and the money is yours.

Suppose you choose door 1. But before Monty Hall opens that door, he opens one of the other two doors, picking one he knows has nothing behind it. Suppose he opens door 2. This takes 2 out of the running, so the only question now is about door 1 vs. door 3.

Monty then lets you reconsider your earlier choice: you can either stick with door 1 or switch to door 3. What should you do? More generally, whatever door a person originally picks, should they switch when given the option?

You are better off switching. The basic idea is that when Monte opens one of the two doors you have acquired new information. To represent things pictorially, we begin by figuring out the various things that could occur in this game. You could pick any of the three doors, and the money could be behind any of the three. We represent this situation in Figure 27.12 (ignore the thumbs-up signs for now).

The first three numbers on the left represent the door you select (either 1, 2, or 3). And for each selection you could make, there are three places where the money could be. So there are nine possibilities. In the figure the arrow points to the situation where you selected door 2 but the money was behind door 3.




Figure 27.12: Tree of Possible Outcomes in Monte Hall Problem

**How to Read the Table** The rows involving switching doors and its outcomes are shaded. Still, Table 27.1 on the following page contains a good deal of information so let's consider a couple of examples.

\$	Original pick	Monte opens	Switch to	Result	Stay	Result
1	1	2 or 3	3 or 2	L	1	W
1	2	3	1	W	2	L
1	3	2	1	W	3	L
2	1	3	2	W	1	L
2	2	1 or 3	3 or 1	L	2	W
2	3	1	2	W	3	L
3	1	2	2	W	1	L
3	2	1	3	W	2	L
3	3	1 or 2	2 or 1	L	3	W

Table 27.1: Using a Picture to Solve the Monte Hall Problem

- Row one: I picked door 1 and that's where the money is. Monty can open either door 2 or door 3, since the money isn't behind either. The only place I can move to is the door he didn't open. If he opened door 2 I can move to 3; if he opened 3 I can move to 2. In this case, if I switch I lose.
- Row two: I picked door 2 and the money is behind door 1. The only door Monty can open is door 3. If I stay with door 2 I lose; if I switch to door 1 I win.

If I always stick with my original pick I will win in only 3 of the 9 cases (rows 1, 5, and 9). If I adopt the policy of switching, I will win in 6 of 9 cases (all the other rows). So I am *twice as likely* to win the money if I switch.

We can present the information more compactly in Table 27.2. Each column indicates what happens in each of the nine situations if I stay or if I switch.

Original pick	1	1	1	2	2	3	3	3	3
Door with \$	1	2	3	1	2	3	1	2	3
Stay	W	L	L	L	W	L	L	L	W
Switch	L	W	W	W	L	W	W	W	L

Table 27.2: Simplified Diagram of Monte Hall Problem



We can also represent these outcomes with a tree; indeed the thumbs-up signs in Figure 27.12 on page 561 indicate the cases where you will win if you switch doors. Like our other representations, they show that you double your chances of winning if you switch.

As usual, it is helpful to think in terms of proportions or frequencies. Imagine that you play the game many times. What should you expect to happen, on average, every one hundred times played? You would win about sixty six of the hundred games if you switch. You would lose sixty six if you stick with your original door.

### Pictorial Solution to the Two Aces Problem

Remove all of the cards except the aces and eights from a deck. This leaves you with an eight-card deck: four aces and four eights. From this deck, deal two cards to a friend. He is honest and will answer exactly one question about his cards. The issue isn't one about which order in which he got his two cards; indeed, let us suppose he shuffled them before looking at them so that he doesn't even know which of his cards was dealt first.

1. You ask if his hand contains at least one ace. He answers Yes. What is the probability that both of his cards are aces?
2. You ask if his hand contains a red ace. He answers yes. What is the probability that both of his cards are aces?
3. You ask if his hand contains the ace of spades. He answers yes. What is the probability that both of his cards are aces?

Using a bit of makeshift but obvious notation we want to know

$$1^* \Pr(2 \text{ aces} \mid \text{at least 1 ace})$$

$$2^* \Pr(2 \text{ aces} \mid \text{at least 1 red ace})$$

$$3^* \Pr(2 \text{ aces} \mid \text{ace of spades})$$

All three probabilities are different. We can see the basic point with fewer cards, but the goal here is to examine a problem that needs a little bit bigger diagram than fewer cards would require.

The most difficult part of this problem is figuring out how to represent the various possible hands. We note that there are eight cards one could get on the first draw and seven on the second so when we take order into account we have  $8 \times 7 = 56$  hands. Order isn't relevant here, however, since we aren't concerned with what card came first, so we can cut 56 in half (counting the outcome of eight

of hearts on the first draw and ace of hearts on the second and the reverse order with the same two cards as the same hand; if you want to consider all 56 cases you can, but it will mean more work).

Indeed, if worst came to worst you could simply write down all of the possible hand in some systematic way and you would end up with the same 28 hands we'll begin with here. This takes work in the present case, but with simpler problems it is entirely feasible.

$A_s A_h$	$A_s A_d$	$A_s A_c$	$A_s E_s$	$A_s E_h$	$A_s E_d$	$A_s E_c$
	$A_h A_d$	$A_h A_c$	$A_h E_s$	$A_h E_h$	$A_h E_d$	$A_h E_c$
		$A_d A_c$	$A_d E_s$	$A_d E_h$	$A_d E_d$	$A_d E_c$
			$A_c E_s$	$A_c E_h$	$A_c E_d$	$A_c E_c$

Table 27.3: Diagram for Thinking about Two Aces Problem

There are more efficient ways to solve this problem than by enumerating the hands, but such an enumeration can sometimes give us a better feel for the situation. Such an enumeration is given in Figure 27.3, and once we have it we can determine the answers simply by inspecting the proportions of various sorts of hands (ones with at least one ace, one with at least one red ace, and so on).

In all three scenarios Wilbur has at least one ace, so only hands with at least one ace are relevant here. A little calculation, or a little trial-and-error doodling, will show that there are 22 hands with at least one ace.

Since the ace of spades showed up in the problem, we can begin by enumerating all of the two-card hands that contain it. Then we can enumerate those with the ace of hearts, diamonds, and clubs, taking care not to list the same hand twice (if worst comes to worst you can simply list all of the hands and then go through and cross out any duplications).

We are dealing with *conditional probabilities* here. Remember how these cut down the set of possibilities (what is sometimes called the *sample space*). You roll a die and can't see where it landed. When Wilma tells you it came up an even number, this eliminates all the odd numbers from consideration and the probability you rolled a two goes up from  $1/6$  to  $1/3$ . Similarly When Wilbur tells you he has at least one ace, this eliminates all of the hands without any aces, and it also changes various probabilities.

**Probability of two aces given at least one** First we'll see what the probability of Wilbur's having two aces is given that he has at least one ace. The 22 hands enumerated in our diagram all contain at least one ace, and 6 of them (those shaded




Figure 27.13: Two Aces Given at Least One

in Figure 27.13) have two aces. So there are 6 ways out of the 22 of having two aces, for a probability of  $6/22$  (which reduces to  $3/11 = .27$ ).

**Probability of two aces given the ace of spades** Now we'll see what the probability of two aces is given that one of the cards is the ace of spades. This cuts down the cases we consider—our sample space—to just those hands with the ace of spades. These occupy the first row of our diagram, and there are exactly seven of them. They are enclosed with a dashed line in Figure 27.14.




Figure 27.14: Two Aces Given the Ace of Spades

Of the seven hands that contain the ace of spades, only three contain two aces (these are enclosed with a solid line and shaded). So there are 3 two-ace hands out of 7. This yields a probability of  $3/7 (= 0.43)$ . Hence the probability of two aces given an ace of spades is substantially higher than the probability of two aces given at least one ace. We leave the last of the three puzzles as an exercise.

### Decision Trees

We can often represent situations where we have to make a difficult decision by trees. In cases where we have a good estimate of the relevant probabilities and payoffs, we can also represent these on the tree and use it to help us calculate expected values (p. 308). Here we will simply note the way in which trees, even in the absence of such information, can help us sort of the relevant alternatives and help us focus those that matter the most.

**Decision trees:** picturing the contingencies

We will consider a hypothetical example, but situations like this are common; at some point virtually all of us have to make difficult, very possibly life-or-death, decisions about medical treatment, either for ourself, or a child or an aging parent.

John's test results are not good. They suggest some possibility of a brain tumor, although the probability is considerably less than 50%. If it *is* cancer and it's not surgically removed John will die in the next year or so. But there is also a genuine risk from the surgery; in a few cases it is fatal, and it leads to lasting damage in other cases. Unless John has the surgery within two months, it will be too late, so the decision needs to be made soon. Of course John could choose just to not think about it, but this amounts to the decision not to have surgery.




Figure 27.15: Simple Decision Tree

We can represent John's situation with a tree like that in Figure 27.15. This helps us chart the various possibilities. In real-life we rarely can obtain very precise values for probabilities, but in some medical situations a good deal is known about the percentage or proportion of those who have cancer given certain test results or who die in surgery of a given sort. Where we have these, even if they aren't completely precise, we can incorporate them into our tree. Such trees can be developed in considerable detail, but we won't explore them further here.

### Rectangular Diagrams for Logic and Probability

In this section we develop a simple method for representing important logical and probabilistic relationships with 2 by 2 diagrams. It takes a few minutes practice to become skilled with these diagrams, but it's a good investment, because it allows us to see many important logical and probabilistic relationships.<sup>1</sup>

Where there are only two variables or factors, say  $A$  and  $B$  it is often very useful to represent them with a  $2 \times 2$  rectangle like that in Figure 27.16.

**Columns** represent  $A$  and its negation  $\neg A$ .

**Rows** represent  $B$  and its negation  $\neg B$ .

For easy reference we can think of these four areas as numbered regions, starting with 1 in the lower right and working our way around (Figure 27.16).

	$A$	$\neg A$
$B$	3	2
$\neg B$	4	1

Figure 27.16: Numbering the Regions




Figure 27.17: Areas Corresponding to  $A$  and  $B$

In the diagram on the left in Figure 27.17 the region with horizontal (left-to-right) hatching (quadrants 3 and 4) represents the area of the rectangle in which  $A$  is true. The area with vertical (up-and-down) hatching (quadrants 1 and 2) represents the remaining region where  $A$  is false, so the vertical region represents  $\neg A$ . Similarly, the horizontal region of the diagram on the right in Figure 27.17 represents the area where  $B$  is true and the vertical region represents that in which  $B$  is false.

<sup>1</sup>This section can be skipped without loss of continuity. If you do skip it, you should go to the Chapter Exercises beginning on p. 570.




Figure 27.18: Regions Corresponding to Conjunction and Disjunction

When we are thinking about probability we interpret the total area of the rectangle as 1. This shows us why the probability of a negation like  $\neg A$  will be whatever portion of the unit of probability that isn't taken by  $A$ .  $A$  and  $\neg A$  have to divide the one unit of probability between them, so  $\Pr(\neg A) = 1 - \Pr(A)$ .

In the lefthand diagram of Figure 27.18 the white, unshaded region (quadrant 3) is where both  $A$  and  $B$  are true, so it represents the situation where the *conjunction*,  $A \& B$ , is true. The remaining, shaded regions (1, 2, and 4) are where this conjunction is false, which is just to say the region where the conjunction's negation,  $\neg(A \& B)$ , is true.

In the right subfigure in Figure 27.18 we see that the *disjunction*,  $A \text{ or } B$ , is true in all of the regions except region 1. The white area represents the area where the disjunction is true, so the shaded area represents the area where it is false, i.e., the area (just quadrant 1) where its negation,  $\neg(A \text{ or } B)$  is true.

But we see that this is also the exact area in which  $A$  is false and  $B$  is false. After all,  $A$  is false in the right hand column and  $B$  is false on the bottom row. And this column and row overlap in the first quadrant. So this bottom right region is the area where  $\neg A \& \neg B$  is true.

Since the regions for  $\neg(A \text{ or } B)$  and  $\neg A \& \neg B$  coincide (righthand diagram in Figure 27.18), they say the same thing. This is a one of two equivalences known as *De Morgan's Laws*. In terms our of regions:

$\neg(A \text{ or } B) = \neg A \& \neg B$ . The claim that neither  $A$  nor  $B$  is true is equivalent to the claim that both  $A$  and  $B$  are false.

There is another, mirror image, version of De Morgan's laws:

$\neg(A \& B) = \neg A \text{ or } \neg B$ . The claim that  $A$  and  $B$  are not both true is equivalent to the claim that at least one or the other is false.

We relabel an earlier diagram as Figure 27.19 on the next page to illustrate this. Here the non-shaded region is  $A \& B$  so the shaded region is its negation,  $\neg(A \& B)$  (in words: not both  $A$  and  $B$ ). But with a bit of looking you can see that this cor-

responds to the region in which  $A$  is false or  $B$  is false (or both), i.e., to  $\neg A$  or  $\neg B$




Figure 27.19: Shaded:  $\neg(A \& B) = \neg A$  or  $\neg B$

Equivalence here amounts to two-way validity, so our diagrams show that four separate argument patterns are valid (Figure 27.20).

$$\frac{\neg(A \& B)}{\text{So } \neg A \text{ or } \neg B} \quad \frac{\neg A \text{ or } \neg B}{\text{So } \neg(A \& B)} \quad \frac{\neg(A \text{ or } B)}{\text{So } \neg A \& \neg B} \quad \frac{\neg A \& \neg B}{\text{So } \neg(A \text{ or } B)}$$

Figure 27.20: Four Valid Argument Patterns

We apply all this quite directly to probability by taking the total area of a two-by-two rectangle to have the area 1 (representing the total amount of probability). Sentences represented by exactly the same area must have exactly the same probability. Hence, by De Morgan’s Laws,

$$\Pr(\neg(A \text{ or } B)) = \Pr(\neg A \& \neg B) \quad \text{and} \quad \Pr(\neg(A \& B)) = \Pr(\neg A \text{ or } \neg B).$$

We can also use rectangular diagrams to illustrate several of our rules for calculating probabilities. Recall the rule for disjunctions with incompatible disjuncts. To say that they are incompatible is to say that there is *no overlap* in their areas of the diagram.




Figure 27.21: Rectangular Diagrams for Disjunctions

In the diagram on the left in Figure 27.21  $A$  is represented by the horizontal hatching and  $B$  by the vertical hatching. Since  $A$  and  $B$  do not overlap (quadrant

3 is empty), we simply add their probabilities to get to probability for  $A$  or  $B$ . By contrast in the diagram on the right in Figure 27.21 on the preceding page  $A$  and  $B$  do overlap and so when we add their areas we add  $A \& B$  (quadrant 3) twice. To make up for this we have to subtract it once.  $\Pr(A \text{ or } B) = \Pr(A) + \Pr(B) - \Pr(A \& B)$ .

## 27.4 Chapter Exercises

1. Draw a diagram like those in Figure 27.9 on page 558 that indicates a stronger negative correlation between smoking and heart attacks.
2. Your dog has a litter of four. Which proportion of males and females is more likely? (a) two of each, (b) three of one and one other the other (i.e., three males and one female or else three females and one male). Draw a diagram to determine the answer. (From “Ask Marilyn,” *Parade Magazine*, August 10, 1997.)
3. Inspect Figure 27.3 on page 564 to determine the probability of two aces given at least one red ace.
4. Go back to the two aces problem and write out the remaining hands, those that contain only eights.
5. If you (or someone you cared a lot about) was in John’s position (Figure 27.15 on page 566), what additional information would you want? How would it help you make your decision? What would you do if you were John?<sup>2</sup>

---

<sup>2</sup>Jill. H. Larkin & Herbert A. Simon, “Why a Diagram is (Sometimes) Worth Ten Thousand Words,” *Cognitive Science*, 11, 1987; 65–99 is an excellent discussion of the usefulness of diagrams. The card selection task was invented by the psychologist P. C. Wason, “Reasoning about a Rule,” *Quarterly Journal of Experimental Psychology*, 20, 1968; 273–281. For an excellent extended treatment of diagrammatic reasoning in logic see Jon Barwise and John Etchemendy, *Language, Proof, and Logic*. Stanford: CSLI Publications, 1999. Further references to be supplied



## 27.5 Appendix: Inverse Probabilities and Bayes' Rule

Only two cab companies operate in Belleville. The Blue Company has blue cabs, and the Green Company has Green Cabs. Exactly 85% of the cabs are blue and 15% are green. A cab was involved in a hit-and-run accident at night. An eyewitness, Wilbur, identified the cab as a Green cab. Careful tests were done to ascertain peoples' ability to distinguish between blue and green cabs at night. The tests showed that people were able to identify the color correctly 80% of the time, but they were wrong 20% of the time. What is the probability that the cab involved in the accident was indeed a green cab, as Wilbur says?<sup>3</sup>

Problems like this require us to *integrate* knowledge about the present case (here, what the eyewitness says) with prior information about base rates (here, what proportion of the cabs are green) In many cases we focus too much on the present information and ignore information about base rates. The correct way to give both pieces of information their proper due is to use Bayes' Rule.

### Bayes' Rule

Bayes' Rule (named after the Reverend Thomas Bayes, who discovered it in the eighteenth century) is really just another rule for calculating probabilities. It tells us how we should modify or update probabilities when we acquire new information. It gives the *posterior probability* of a hypothesis  $A$  in light of a piece of new evidence  $B$  as

$$(1) \quad \Pr(A|B) = \frac{\Pr(A) \times \Pr(B|A)}{\Pr(B)}$$

We say that  $\Pr(A)$  is the *prior probability* of the hypothesis,  $\Pr(B|A)$  the *likelihood* of  $B$  given  $A$ , and  $\Pr(B)$  the *prior probability* of evidence  $B$ . For example,  $A$  might be the hypothesis that Smith is infected by the HIV virus and  $B$  might be the datum (piece of evidence) that his test came back positive.

By way of illustration let's work through example one above (we will work through example two below). This example involved the Blue Company, which has all blue cabs, and the Green Company, which has all green cabs. Exactly 85% of the cabs are blue and the other 15% are green. A cab was involved in a hit-and-run accident at night. An honest eyewitness, Wilbur, identified the cab as a green cab. Careful tests were done to ascertain witness' ability to distinguish between

---

<sup>3</sup>This appendix can be skipped without loss of continuity; it is more difficult than the rest of the material in this book.

blue and green cabs at night; these showed that people were able to identify the color correctly 80% of the time, but they were wrong 20% of the time.

What is the probability that the cab involved in the accident was indeed a green cab? The fact that witnesses are pretty reliable leads many people to suppose that Wilbur probably right, even that the probability that he is right is .8. But is this correct?

A large part of the battle is often setting up clear notation, so let's begin with that:

- $\Pr(G) = .15$  This is the *base rate* of green cabs in the city. It gives the *prior probability* that the cab in the accident is green. Similarly,  $\Pr(B) = .85$ .
- $\Pr(SG|G) = .80$  This is the probability the witness will be correct in saying green when the cab in fact is green, i.e., given that the cab really was green. Similarly,  $\Pr(SB|B) = .80$ .

These are the probabilities that witnesses are correct, so by the negation rule, the probabilities of misidentifications are:

- $\Pr(SG|B) = .2$  and  $\Pr(SB|G) = .2$

What we want to know is the probability that the cab really was green, given that Wilbur said it was green, i.e., we want to know  $\Pr(G|SG)$ .

According to Bayes' Rule, this probability is given by:

$$\Pr(G|SG) = \frac{\Pr(G) \times \Pr(SG|G)}{\Pr(SG)}$$

We have the values for the two expressions in the numerator— $\Pr(G) = .15$  and  $\Pr(SG|G) = .8$ —but we have to do a little work to determine the value for the expression— $\Pr(SG)$ —in the denominator.

To see what this value—the probability that a witness will say that the cab was green—must be, note that there are two conditions under which Wilbur could say that the cab in the accident was green. He might say this when the cab was green *or* when it was blue. This is a disjunction, so we add these two probabilities. In the first disjunct Wilbur says green *and* the cab is green; in the second disjunct Wilbur says green *and* the cab is blue. Putting this together we get:

$$\Pr(SG) = \Pr(G \& SG) + \Pr(B \& SG)$$

Now our rule for conjunctions tells us that  $\Pr(G \& SG) = \Pr(G) \times \Pr(SG|G)$  and  $\Pr(B \& SG) = \Pr(B) \times \Pr(SG|B)$ . So

$$\begin{aligned}
 \Pr(SG) &= \Pr(G) \times \Pr(SG|G) + \Pr(B) \times \Pr(SG|B) \\
 &= (.15 \times .80) + (.85 \times .20) \\
 &= .12 + .17 \\
 &= .29
 \end{aligned}$$

Finally, we substitute this number, .29, into the denominator of Bayes' Rule:

$$\begin{aligned}
 \Pr(G|SG) &= \frac{\Pr(G) \times \Pr(SG|G)}{\Pr(SG)} \\
 &= \frac{.15 \times .80}{.29} \\
 &= .414
 \end{aligned}$$

So the probability that the witness was correct in saying the cab was green is just a bit above .4—definitely less than fifty/fifty—and (by the negation rule) the probability that he is wrong is nearly .6. This is so even though witnesses are pretty reliable. How can this be? The answer is that the high *base rate* of blue cabs and the low *base rate* of green cabs makes it somewhat likely that the witness was wrong in this case.

### Rule for Total Probabilities

In calculating the value of the denominator in the above problem we made use of the *Rule for Total Probabilities*. We will use a fairly simple, but widely applicable version of this rule here. If a sentence  $B$  is true, then either  $B$  and  $A$  are true or else  $B$  and  $\neg A$  are true (since any sentence  $A$  is either true or false). This allows us to express the probability of  $B$ — $\Pr(B)$ —in a more complicated way.

This may seem like a strange thing to do, but it turns out that it is often possible to calculate the probability of the more complicated expression when it isn't possible to obtain the probability of  $B$  more directly.

$$\begin{aligned}
 \Pr(B) &= \Pr[(B \& A) \text{ or } (B \& \neg A)] \\
 &= \Pr(B \& A) + \Pr(B \& \neg A) \\
 &= [\Pr(A) \times \Pr(B|A)] + [\Pr(\neg A) \times \Pr(B|\neg A)]
 \end{aligned}$$

In short

$$\Pr(B) = [\Pr(A) \times \Pr(B|A)] + [\Pr(\neg A) \times \Pr(B|\neg A)].$$

This rule can be useful in cases that do not involve Bayes' Rule, as well as in many cases that do. It is particularly useful when we deal with an outcome that can occur in either of two ways.

*Example:* A factory has two machines that make widgets. Machine  $A$  makes 800 per day and 1% of them are defective. The other machine, call it  $\neg A$ , makes 200 a day and 2% are defective. What is the probability that a widget produced by the factory will be defective ( $D$ )? We know the following:

- $\Pr(A) = .8$  – the probability that a widget is made by machine  $A$  is .8 (since this machine makes 800 out of the 1000 produced everyday)
- $\Pr(\neg A) = .2$  – the probability that a widget is made by the other machine  $\neg A$  is .2
- $\Pr(D|A) = .01$  – the probability that a widget is defective given that it was made by machine  $A$ , which turns out 1% defective widgets, is .01p
- $\Pr(D|\neg A) = .02$  – the probability that a widget is defective given that it was made by machine  $\neg A$ , which turns out 2% defective widgets, is .02p

Plugging these numbers into the theorem for total probability we have:

$$\begin{aligned}\Pr(D) &= [\Pr(A) \times \Pr(D|A)] + [\Pr(\neg A) \times \Pr(D|\neg A)] \\ &= [.8 \times .01] + [.2 \times .02] \\ &= 0.012\end{aligned}$$

### Odds Version of Bayes' Rule

Particularly when we are concerned with mutually exclusive and disjoint hypothesis (cancer vs. no cancer), the odds version of Bayes' rule is often the most useful version. It also allows avoid the need for a value for  $\Pr(e)$ , which is often difficult to come by. Let  $H$  be some hypothesis under consideration and  $e$  be some newly acquired piece of evidence (like the result of a medical test) that bears on it. Then

$$(2) \quad \frac{\Pr(H|e)}{\Pr(\neg H|e)} = \frac{\Pr(H)}{\Pr(\neg H)} \frac{\Pr(e|H)}{\Pr(e|\neg H)}$$

The expression on the left of (2) gives the posterior odds of the two hypotheses,  $\Pr(H)/\Pr(\neg H)$  gives their prior odds, and  $\Pr(e|H)/\Pr(e|\neg H)$  is the likelihood (or diagnostic) ratio. In integrating information our background knowledge of base rates is reflected in the prior odds, and our knowledge about the specific case is represented by the likelihood ratio. For example, if  $e$  represents a positive result in a test for the presence of the HIV virus, then  $\Pr(e|H)$  gives the *hit rate (sensitivity)* of the test and  $\Pr(e|\neg H)$  gives the *false alarm rate*. As the multiplicative relation

in (2) indicates, when the prior odds are quite low even a relatively high likelihood ratio won't raise the posterior odds dramatically.

The importance of integrating base rates into our inferences is vividly illustrated by example two above. Suppose that we have a test for the HIV virus. The probability that a person who really has the virus will test positive is .90 while the probability that a person who does *not* have it will test positive is .20. Finally, suppose that the probability that a person in the general population has the virus (the base rate of the virus) is .01. How likely is Smith, whose test came back positive, to be infected?

Because the test is tolerably accurate, many people suppose that the chances are quite high that Smith is infected, perhaps even as high as 90%. Indeed, various studies have shown that many physicians even suppose this. It has been found that people tend to confuse probabilities and their converses. The probability that we'll get the positive test,  $e$ , if the person has the virus is .9, i.e.,  $\Pr(e|H) = .9$ . But we want to know the probability of the converse, the probability that the person has the virus given that they tested positive, i.e.,  $\Pr(H|e)$ . These need not be the same, or even close to the same. In fact, when we plug the numbers into Bayes' Rule, we find that the probability that someone in this situation is infected is actually quite low.

To see this we simply plug our numbers into the odds version of Bayes' rule. We have been told that

- $\Pr(e|H) = .9$
- $\Pr(e|\neg H) = .2$
- $\Pr(H) = .01$
- $\Pr(\neg H) = .99$ .

Hence

$$\frac{\Pr(H|e)}{\Pr(\neg H|e)} = \frac{\Pr(H)}{\Pr(\neg H)} \times \frac{\Pr(e|H)}{\Pr(e|\neg H)} = \frac{.9}{.2} \times \frac{.01}{.99} = \frac{1}{22}$$

This gives the odds: 22 to 1 against Smith's being infected. So  $\Pr(H|e) = \frac{1}{23}$ . Although the test is fairly accurate, the low base rate means that there is only 1 chance in 23 that Smith has the virus.

Similar points apply to other medical tests and to drug tests, as long as the base rate of the condition being tested for is low. It is not difficult to see how policy makers, or the public to which they respond, can make very inaccurate assessments of probabilities, and hence poor decisions about risks and remedies, if they overlook the importance of base rates.

As we noted earlier, it will help you to think about these matters intuitively if you try to rephrase probability problems in terms of *frequencies* or *proportions* whenever possible.

### Further Matters

Bayes' Rule makes it clear when a conditional probability will be equal to its converse. When we look at

$$\Pr(A|B) = \frac{\Pr(A) \times \Pr(B|A)}{\Pr(B)}$$

We see that  $\Pr(A|B) = \Pr(B|A)$  exactly when  $\Pr(A) = \Pr(B)$ , so that they cancel out (assuming, as always, that we do not divide by zero). Dividing both sides of Bayes' Rule by  $\Pr(B|A)$  gives us the following ratio

$$\frac{\Pr(A|B)}{\Pr(B|A)} = \frac{\Pr(A)}{\Pr(B)}$$

which is often useful.

**Updating via Conditionalization** One way to update or revise our beliefs in light of the new evidence  $e$  is to set our new probability (once we learn about the evidence) as:

$$\text{New Pr}(H) = \text{Old Pr}(H|e),$$

with  $\text{Old Pr}(H|e)$  determined by Bayes' Rule. Such updating is said to involve *Bayesian conditionalization*.

### The Problem of the Reference Class

There are genuine difficulties in selecting the right group to consider when we use base rates. This is known as the *problem of the reference class*. Suppose that we are considering the effects of smoking on lung cancer. *Which* reference class should we consider: all people, just smokers, just heavy smokers?

Similar problems arise when we think about regression to the mean. We know that extreme performances and outcomes tend to be followed by ones that are more average (that regress to the mean). If Wilma hits 86% of her shots in a basketball game, she is likely to hit a lower percentage the next time out. People do not seem to have an intuitive understanding of this phenomenon, and the difficulty is compounded by the fact that, as with base rates, there is a problem about reference class. *Which* average will scores regress to: Wilma's average, her recent average, her team's average?

Up to a point smaller reference classes are likely to underwrite better predictions, and we will also be interested in reference classes that seem causally relevant to the characteristic we care about. But if a reference class becomes too small, it

won't generate stable frequencies, and it will often be harder to find data on smaller classes that is tailored to our particular concerns.

There is no one right answer about which reference class is the correct one. It often requires a judicious balancing of tradeoffs. With base rates it may be a matter of weighing the costs of a wrong prediction against the payoffs of extreme accuracy, or the value of more precise information against the costs of acquiring it or (if you are a policy maker) the fact that the number of reference classes grows exponentially with each new predictor variable. But while there is no uniquely correct way to use base rates, it doesn't follow that it is fine to ignore them; failing to take base-rate information into account when we have it will often lead to bad policy.

### Exercises on Bayes' Theorem and Conditional Probabilities

1. In the section on the Rule for Total Probability (p. 573) we encountered the factory with two machines that make widgets. Machine  $A$  makes 800 per day and 1% of them are defective. The other machine, call it  $\neg A$ , makes 200 a day and 2% are defective. What is the probability that a widget is produced by machine  $A$ , given that it is defective? We know the probability that it is defective if it is produced by  $A$ — $\Pr(D|A) = .01$ —but this asks for the opposite or *converse probability*; what is  $\Pr(A|D)$ ? Use Bayes' Rule to calculate the answer.
2. Earlier you were asked to draw a picture to solve the following problem; now solve it by calculation and see if your answers agree. Officials at the suicide prevention center know that 2% of all people who phone their hot line actually attempt suicide. A psychologist has devised a quick and simple verbal test to help identify those callers who will actually attempt suicide. She found that
  1. 80% of the people who will attempt suicide have a positive score on this test.
  2. but only 5% of those who will not attempt suicide have a positive score on this test.

If you get a positive identification from a caller on this test, what is the probability that he would actually attempt suicide?

3. A clinical test, designed to diagnose a specific illness, comes out positive for a certain patient. We are told that:
  1. The test is 79% accurate: the chances that you have the illness if it says you do is 79% and the chances that you do not have the illness if it says you don't is also 79%.
  2. This illness affects 1% of the population in the same age group as the patient.

Taking these two facts into account, and assuming that you know nothing about the patient's symptoms or signs, what is the probability that this particular patient actually has the illness?

4. Suppose that you handed two bags of poker chips, but from the outside you can't tell which is which.

1. *Bag 1*: 70 red chips and 30 blue chips
2. *Bag 2*: 30 red chips and 70 blue chips.

You pick one of the two bags at random and draw a chip from it. The chip is red. You replace it, draw again, and get another chip, and so on through twelve trials (i.e., twelve draws). In the twelve trials you get 8 red chips and 4 blue chips. What is the probability that you have been drawing chips from *Bag One* (with 70 red and 30 blue) rather than from *Bag Two* (with 30 red and 70 blue)? Most people answer that the probability that you have been drawing from *Bag One* is around .75. In fact, as Bayes' Rule will show, it is .97. But people often revise their probabilities *less* that Bayes' Rule says they should. People are *conservative* when it comes to updating their probabilities in the light of new evidence. Use Bayes' Rule to show that .97 is the correct value.

5. Wilbur has two children. We run into him at the Mall with a teenage boy he introduces as his son. What is the probability that Wilbur's other child is a boy? Second scenario: Wilbur introduces the boy as his oldest son. Now what is the probability that his other child is a boy? (Hint: think about the two-aces problem above).

6. One thousand people, including you, bought one ticket in the local lottery. There were ten winners, all of whom were notified correctly that they had won. But because of a clerical error, 1% of the people who didn't win also received notifications that they did. You received a letter saying you were a winner. What is the probability that you really did win?

7. *Monte Hall Problem* In an earlier chapter we approached the Monty Hall Problem in an intuitive way. Now we will now verify our earlier answers using Bayes' Rule. Recall that in this problem you imagine that you are a contestant in a game show and there are three doors in front of you. There is nothing worth having behind two of them, but there is \$100,000 behind the third. If you pick the correct door, the money is yours.

You choose A. But before the host, Monte Hall, shows you what is behind that door, he opens one of the other two doors, picking one he knows has nothing behind it. Suppose he opens door B. This takes B out of the running, so the only



question now is about door A vs. door C. Monte ('M', for short) now allows you to reconsider your earlier choice: you can either stick with door A or switch to door C. Should you switch?

1. What is the probability that the money is behind door A?
2. What is the probability that the money is behind door C?

Since we are asking these questions once Monte has opened door B, they are equivalent to asking about:

### Answers to Selected Exercises

1. We are given all of the relevant numbers except that for  $\Pr(D)$ , but we calculated this in the section on total probabilities. By Bayes' Rule:

$$\begin{aligned}\Pr(A|D) &= \frac{\Pr(A) \times \Pr(D|A)}{\Pr(D)} \\ &= \frac{.8 \times .01}{.012} \\ &= \frac{2}{3}\end{aligned}$$

7. Lets give the doors letter names so they don't get confused with out numbers. Then, to solve the Monty Hall problem we must calculate:

- 1'  $\Pr(\$ \text{ behind A} | \text{M opened B})$
- 2'  $\Pr(\$ \text{ behind C} | \text{M opened B})$

Bayes' Rule tells us that  $\Pr(\text{Money behind A} | \text{M opened B}) =$

$$\frac{\Pr(\$ \text{ behind A}) \times \Pr(\text{M opened B} | \$ \text{ behind A})}{\Pr(\text{M opens B})}$$

We know the values for the two items in the numerator:

1.  $\Pr(\$ \text{ behind A})$ : the prior probability that the money is behind door A is  $\frac{1}{3}$  (going in, it could equally well be behind any of the three doors).
2.  $\Pr(\text{M opened B} | \$ \text{ behind A})$  is  $\frac{1}{2}$  (there is a fifty/fifty chance that he would open either door B *or* door C, when the money is behind door A).
3. But  $\Pr(\text{M opens B})$ , the number in the denominator, requires some work.

To see the value of the denominator— $\Pr(\text{M opens B})$ —note that Monte will *never* open door B if the money is behind B. Hence, he could open B under exactly two

conditions. He could open B when the money is behind A or he could open B when the money is behind C. So

$$\begin{aligned}
 \Pr(\text{M opens B}) &= [\Pr(\$ \text{ behind A}) \times \Pr(\text{M opens B} \mid \$ \text{ behind A})] \\
 &+ [\Pr(\$ \text{ behind C}) \times \Pr(\text{M opens B} \mid \$ \text{ behind C})] \\
 &= [1/3 \times 1/2] + [1/3 \times 1] \\
 &= [1/6] + [1/3] \\
 &= 3/6 \\
 &= 1/2
 \end{aligned}$$

Plugging this into the denominator in Bayes' Rule, we get:

$$\begin{aligned}
 \Pr(\$ \text{ behind A} \mid \text{M opened B}) &= \frac{\Pr(\$ \text{ behind A}) \times \Pr(\text{M opened B} \mid \$ \text{ behind A})}{\Pr(\text{M opens B})} \\
 &= \frac{1/3 \times 1/2}{1/2} \\
 &= \frac{1/6}{1/2} \\
 &= 1/3
 \end{aligned}$$

So the probability that the money is behind your door, door A, given that Monte has opened door B, is  $\frac{1}{3}$ . Since the only other door left is door C, the probability that the money is there, given that Monte opened door B, should be  $\frac{2}{3}$ . Use Bayes' Rule to prove that this is correct.

### Derivation of Bayes' Rule

Where does Bayes' theorem come from? It is straightforward to derive it from our rules for calculating probabilities. You don't need to worry about the derivation, but here it is for anyone who likes such things.

$$\begin{aligned}
 \Pr(A \mid B) &= \frac{\Pr(A \& B)}{\Pr(B)} \\
 &= \frac{\Pr(A) \times \Pr(B \mid A)}{\Pr(B)}
 \end{aligned}$$

In actual applications we often don't have direct knowledge of the value of the denominator,  $\Pr(B)$ , but in many cases we do have enough information to calculate it using the Rule for Total Probability. This tells us that:

$$\begin{aligned}
 \Pr(B) &= \Pr(B \& A) + \Pr(B \& \neg A) \\
 &= [\Pr(A) \times \Pr(B \mid A)] + [\Pr(\neg A) \times \Pr(B \mid \neg A)]
 \end{aligned}$$

So the most useful version of Bayes' Rule is often

$$\Pr(A|B) = \frac{\Pr(A) \times \Pr(B|A)}{[\Pr(A) \times \Pr(B|A)] + [\Pr(\neg A) \times \Pr(B|\neg A)]}$$

Bayes' Rule can take more complex forms, but this is as far as we will take it in this book.<sup>4</sup>

---

<sup>4</sup>Many of the ideas and experimental results described in this appendix and in earlier chapters come from an extremely important series of papers by Amos Tversky and Daniel Kahneman. Many of these are collected in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under Uncertainty: Heuristics and Biases*, Cambridge University Press, 1982. Richard E. Nisbett and Lee Ross's *Human Inference: Strategies and Shortcomings of Social Judgment*, Englewood Cliffs, NJ: Prentice-Hall, 1980, contains a detailed and accessible discussion of many of the topics of this chapter. D. M. Eddy, "Probabilistic Reasoning in Clinical Medicine: Problems and Opportunities," in the book edited by Kahneman, Slovic & Tversky documents our strong tendency to confuse conditional probabilities with their converses (also called their *inverses*). The two king problem is presented by Irving Copi in his *Introduction to Logic*, 3rd/ed., Macmillan, 1968. The problem about the two sons, along with many other interesting examples, may be found in Maya Bar-Hillel and Ruma Falk's "Some Teasers Concerning Conditional Probabilities," *Cognition*, 11 (1982) 109-122.



## Chapter 28

# Recognizing where Cognitive Tools Apply: Cues, Transfer, and Habits

**Overview:** You will only be able to apply the concepts and skills you learn here in situations outside the classroom if you can you *recognize* those situations as ones where an appropriate cognitive tool is relevant. This is often difficult, and in this chapter we will examine ways to become more sensitive to *cues* that signal the applicability of particular concepts and skills. We also discuss three very general slogans that, if frequently called to mind, would improve a good deal of reasoning: Consider Alternatives! Is there Invisible Data? Is some data *too* visible? We conclude with a discussion of transfer of reasoning skills to situations outside the classroom and the steps for developing the habits necessary for transfer to occur.

### Contents

---

<b>28.1 Good Thinking</b> . . . . .	<b>584</b>
<b>28.2 Feedback: Learning from Experience</b> . . . . .	<b>588</b>
<b>28.3 Recognizing the Relevance of a Cognitive Tool</b> . . . . .	<b>589</b>
28.3.1 Recognizing when these Tools are Relevant . . . . .	590
28.3.2 Cues That Signal when a Tool Applies . . . . .	592
<b>28.4 Consider Alternatives</b> . . . . .	<b>593</b>
28.4.1 Is there “Invisible Data”? . . . . .	599

---

28.5 Acquiring Cognitive Skills . . . . .	601
28.6 Chapter Exercises . . . . .	604

---

## 28.1 Good Thinking

The world is changing rapidly, and much what you learn now will be outdated in a few years. Many of your grandparents, and perhaps even your parents, had just one or two jobs during their adult life, but studies show that many people now in their twenties will have a succession of jobs, often in quite different fields, and this will require the acquisition of new skills.

By contrast, the tools and skills learned for critical reasoning will not date. They are things you can use to *learn how to learn*. But if they are going to help you, you must be able to apply the appropriate concepts and principles to new situations. This is easier said than done, and in this chapter we will focus on ways to make it a little easier.

### What Counts as Good: Who Decides?

Throughout the book we have recommended some ways of reasoning as good and criticized others as flawed, often labeling the latter as fallacies, biases, or errors. But who decides what counts as good or bad reasoning? Indeed, nowadays one sometimes hears that the strategies and tools recommended in books like this foster rigid linear thinking, old-fashioned thinking within the box that limits, rather than improves, reasoning.

Talk like this sounds better in the abstract, though, than it does when we get down to concrete cases (as we will in a moment). The point of reasoning is typically to help us decide what to *do*, to assist us in our actions, and the worst thing about bad reasoning is that it often leads to actions that no one, not even the person who acts, wanted or intended.

Everyone makes mistakes in reasoning that are costly, to themselves and to others. You only have to glance at the morning headlines or the evening news to hear about peoples' miscalculations and follies, about reasoning and actions that are utterly self-defeating, even by the agents' own standards. Indeed, you have probably seen instances of this first hand. Here are three illustrations:

**Bay of Pigs** On April 17, 1961 John F. Kennedy launched an attempt to overthrow Castro in Cuba. The operations was a complete failure, and it cast doubt on

Kennedy's ability to govern. The assumptions and planning of Kennedy and his advisors were full of mistakes that it led to this, the so-called Bay of Pigs, fiasco (p. 489).

**Clinton and Lewinsky** Bill Clinton had the most powerful job in the world, one he dreamed of all his life. But he almost lost it, and a great deal else, because of his irresponsible behavior with Monica Lewinsky.

**Enron** In the fall of 2001 the giant company Enron collapsed. Many people, including employees, had invested a great deal of their money in it. The danger signs probably seem clearer in retrospect (as a result of high-sight bias), but even when things began to *look* shaky, many people continued to invest. Similar problems occurred, on a smaller scale, with the collapse of many dot.com companies.

Many of the pitfalls and problems discussed in earlier chapters affect us in ways that *matter* to us a lot. For example, issues about framing, anchoring, and contrast effects arise in many negotiations: determining who does which parts of the work around the house, buying (or selling) a car or a house, trying to get the best salary you can at the new job, or (more depressingly) deciding who gets what, including access to the children, after a divorce.

Here are just a few of the ways good and bad reasoning have an impact on our lives.

**Framing Effects** Our reasoning and choices about options often depend on the way the options are framed or described. If an issue is framed in terms of a gain we are likely to think about it one way; if it is framed as a loss we are more likely to think about it differently (p. 349). If we aren't on guard, our choices may not be based on the real features of situation, but on the details, often quite trivial ones, of how the situation is described. Furthermore, since other people often supply the frames, this makes us vulnerable to manipulation.

We also noted a study (p. 350) where physicians recommended different treatments (surgery or radiation therapy), depending on how the *same* situations (involving cancer) were described. But no one would want to their cancer treatment determined by something as irrelevant as this.

**Conditional Probabilities** We have a tendency to confuse conditional probabilities with their inverses (i.e., we tend to confuse  $\Pr(A|B)$  with  $\Pr(B|A)$ , p. 554). This can lead to many problems. For example, such probabilities are easy to confuse when we interpret medical tests (e.g., a test for aids), drug tests, and lie de-

tector tests. Indeed, diagnoses and treatment are often based on beliefs about conditional probabilities (e.g., the probability of recovery given treatment *A* vs. the probability of recovery given treatment *B*), beliefs that even some physicians get wrong (p. 575). You don't want mistakes like this to occur when a doctor is treating your child for a serious disease or when a prospective employer learns that you failed a drug test.

**Fallacy of Irrelevant Reason** We commit this fallacy if we base our reasoning or conclusions on irrelevant premises (p. 178). Premises that are irrelevant are ones that simply *do not bear* on truth or falsity of the conclusion one way or another, so they can never provide good reasons for believing a conclusion. When we commit this fallacy we end up believing things *for no (good) reason at all*.

**Cumulative Risk** Many of the things we do are relatively safe each time we do them. For example, the probability that you will be in an automobile accident on any given outing or have a single condom fail are low. But when we do things over and over we are dealing with probabilities of disjunctions. These mount up, so that over a period of time the risk can become reasonably large (p. 315). This is a simple fact about probabilities, and ignoring it leads to bad risk management.

**Availability** If we base our inductive inferences on samples that come readily to mind—that are available (p. 331)—we will often be basing our inductive reasoning on biased samples, and this is likely to lead us to conclusions that are simply false. If we don't appreciate how large the difference between the probability of having a heart attack and the probability of dying at the hands of a terrorist are (even after September 11), it will be difficult to make rational plans about smoking, diet, and travel. Thousands and thousands of Americans die from heart disease every year (in 1997 over 725,000 Americans died this way), whereas even now very few Americans die at the hands of terrorists.

**Causal Fallacies** If our actions are to be effective, we need to act in ways that bring about—that *cause*—the effects we want. If we reason badly about causation (p. 394), so that we don't really know what causes what, many of our predictions will be false and many of our plans and decisions based on those (faulty) predictions will lead to outcomes that we don't desire.

Problems also arise when we confuse necessary and sufficient conditions, commit the fallacies discussed in Module Four, fall victim to confirmation bias or illusory correlations, ignore regression to the mean, underutilize base rates, and so on and on.



In many of these cases our own bad reasoning opens the door for others to manipulate us, e.g., by framing things in a way favorable to them (rather than us), by setting anchors that skew our reasoning, by manipulating our actions and attitudes by dissonance reduction mechanisms (Ch. 19) or by the various tools of professional persuaders (p. 454).

Wanting to avoid being at the mercy of flawed reasoning and manipulation is not just some prejudice in favor of “linear” reasoning. It is an essential part of wanting to be in control of our own life and to base our actions on a realistic assessment of the facts, rather than on irrelevant information or false assumptions.

**How Good, or Bad, are We?** How good—or bad—*is* our reasoning? People who study the matter are not in complete agreement, though in many ways we are stunningly good. No one has a clue how to build a computer with a perceptual system or memory remotely as good (at some things) as ours. And even in cases where we have our lapses, our judgments and decisions are not hopelessly flawed. We couldn’t make it through the day if they were.

The spotty picture emerging from several decades of research suggests that we have the ability to reason pretty well under some conditions. But this ability is fragile and unstable, and it can be affected, diverted, and even subverted by a variety of influences, as the follies and fiascoes noted above make clear.

It doesn’t really matter exactly how good—or bad—we are, though. The important points are that we could all do better and that we would be better off if we did. You don’t need to know whether the rate of a serious disease is 12% or 20% before you try to find ways to cure it. Similarly, we don’t need to know precisely how widespread bad reasoning is in order to try to improve it.

**No Quick Fixes** In many cases the first step to better reasoning is to learn about tempting, but faulty, ways of reasoning like ignoring base rates or regression effects. But, unfortunately, there is a great deal of evidence that merely learning about such pitfalls, and being warned against them, is not an effective way to avoid them, especially over the long run.

It would help us understand how to reason better if we knew more about why we reason badly. We can often help sick people feel better by treating their symptoms, e.g., giving them aspirin to ease the aches and pains. But it is usually more effective if we can identify the underlying cause of an illness (e.g., infection by a certain bacterium) and intervene to change things at the level of root causes (e.g., by killing the bacteria with an antibiotic).

Similarly, it would probably be easier to improve reasoning by designing interventions that work on the underlying mechanisms that produce suboptimal rea-

soning, rather than by trying to work on the symptoms directly (e.g., by nagging at people to use base rates).

We do have some understanding of the various causes of suboptimal reasoning, but we do not yet know nearly as much as we would like. In cases where we do know a little, we can try to change the basic causes of bad reasoning; in cases where we don't, we can still try to treat the symptoms.

Unfortunately, there are no magic bullets for either sort of case. It is more accurate to think of the techniques we have studied in much way we think about exercise and a healthy diet. Neither guarantees that things will always go well—that we will always be healthy or that we will always reason carefully—but with them things will go well more often than they would with a bad diet, or with faulty reasoning. Moreover, again as with staying healthy, learning to reason better isn't an all-or-nothing proposition. It is a life-long process, and all of us have room for improvement.

In a moment we will examine some general ways we can improve our reasoning, but first lets ask how each one of us—individually—might discover how good, or how flawed, our *own* reasoning is.

## 28.2 Feedback: Learning from Experience

Ironically, suboptimal reasoning is one obstacle to finding out how well we really *do* reason. We learn many of the things we do by trial and error: we often don't get things right the first time, but frequently we can improve if we see what went wrong and try to adjust to do better. This only works, though, *if* we note the errors; without **feedback** telling us how accurate our reasoning has been, we can't learn from our mistakes.

There are various reasons why it can be difficult to get good feedback about the quality of our reasoning, but the common pitfalls in reasoning we have encountered in earlier chapters are among the culprits. Here are a three quick examples.

**Hindsight bias** Hindsight bias is the tendency to overestimate the likelihood that we would have predicted an outcome *after* we learn that it occurred (p. 145). It impedes accurate feedback, because it confirms our view that we are right (“I knew what would happen”) more often than we actually are. This makes it more difficult to correct our mistakes by learning from past errors, since we underestimate how often we are in error.

**Confirmation Bias and Availability** Since most of us like to think we are reasonably accurate, we may notice and remember evidence showing our success,

while ignoring evidence that suggests failures. This positive evidence will then be more available in memory, so we will be inclined to think that we have a better track record than we actually do.

**Feedback can be Unpleasant** Feedback can be painful. Wishful thinking, various self-serving biases, and perhaps even defense mechanisms like denial can lead us to avoid feedback. After all, we may find out that we didn't do as well as we had thought or hoped—maybe we didn't do very well at all. I once had a colleague who found reading his teaching evaluations so unpleasant that he simply didn't read them. He probably could have learned things that would have improved his teaching (and led to better evaluations). But he preferred not to know.

Since these obstacles to getting good feedback involve flawed reasoning, we would get more accurate feedback if we did a better job of reasoning.

## 28.3 Recognizing the Relevance of a Cognitive Tool

### Cognitive Tools

The emphasis throughout this book has been on learning about, understanding, and *using* a diversified toolkit of *cognitive tools*. These are concepts and principles and rules that can help us reason more carefully and accurately.

**1. Error Avoidance** Some of these tools are used primarily for spotting and avoiding faulty reasoning. Many of these come with the labels of *fallacy*, *bias*, or *error*. For example, we learned about the fallacy of affirming the consequent, the strawman fallacy, and the gambler's fallacy. We also learned about self-serving biases, confirmation bias, and hindsight bias, and we encountered the fundamental attribution error. And we studied strategies for avoiding errors in perception and memory, recognizing regression to the mean, and being more sensitive to base rates. These tools for recognizing problems are important, because the things they help us avoid are sure-fire ways to reason badly.

**2. Positive Tools** We also learned to use tools for reasoning well. For example, we learned about valid arguments, necessary and sufficient conditions, inductive strength, rules of determining probabilities, evaluating samples, detecting correlations, tracking down causes, assessing risk, and drawing diagrams to make reasoning easier.

**3. Getting Good Information** We can't reason without something to reason *about*, and we also spent some time on tools to help us find, and evaluate, information from various sources.

The three sorts of tools shade off into one another. The key point here is that the more *skilled* you become in the use of these tools, the better your reasoning will be.

### 28.3.1 Recognizing when these Tools are Relevant

The main problem: not realizing when a particular cognitive tool is applicable

The greatest obstacle to reasoning well in the real world does not involve mastering the fine points of concepts and principles and other cognitive tools. The greatest problem is *recognizing situations where they apply*. When the times comes where we really could benefit from a given tool, we're often oblivious; it simply doesn't register that a given tool would help us in the situation we are in. You can learn all the checklists and definitions and rules in the world, but if you can't *recognize* situations where they apply, they are useless. Finding a way to overcome this obstacle is the greatest challenge in teaching, *and learning*, critical reasoning.

#### Recognition is the Key

The general point that recognition is the key will be clearer if we consider several examples.

**Example 1: Sample Size** Earlier we encountered the two hospitals in Smudsville (p. 284). About 60 babies are born everyday in the larger one, and about 15 are born everyday in the smaller one. On average 50% of the births in both hospitals are girls and 50% are boys, but the number varies a some from day to day. Which hospital, if either, is more likely to have more days per year when over 70% of the babies born are boys?

Many people say that the number of 70%-boy days should be about the same in each hospital, but in fact it's more likely in the smaller one. Since about half of all births are boys and about half are girls, the true percentages in the general population are about half and half. Since a smaller sample will be less likely to reflect these true proportions, the smaller hospital is more likely to have more days per year when over 70% of the births are boys. The births at the smaller hospital constitute a smaller sample.

We all know that bigger samples are likely to be more representative of their parent populations, but we sometimes fail to realize that we are dealing with a problem that involves samples. We don't recognize or identify it *as* a problem involving sample size. We don't represent it to ourselves as ones where our tools

for thinking about samples and populations are relevant. Once we recognize that the problem involves sample size, most of us see what the answer must be.

**Example 2: Probabilities of Conjunctions** Consider a frequency version of the conjunction fallacy: Are there are more six letter words ending in ‘ing’ than having ‘n’ as their fifth letter? All six letter words ending in ‘ing’ have ‘n’ in fifth place—and ‘i’ in forth and ‘g’ in sixth. So if we suppose there are more six-letter ‘ing’ words, we commit the conjunction fallacy (p. 314).

But when people get this wrong they are rarely thinking of it as a conjunction and then making an error in reasoning about probabilities of conjunctions. They simply don’t think of it as a conjunction in the first place. They don’t *recognize* or *code* it as a conjunction, and so their tools for thinking about probabilities of conjunction are never applied. When people do see it as a conjunction, especially if they draw a diagram, they tend to get it right.

**Example 3: Anchoring and Adjustment** We fall victim to an anchoring effect whenever some number or other reference points leads us to skew our reasoning in the direction of the anchor (p. 340). We all know that we should ask whether an anchor is a reasonable one or not, but the problem is not that we typically identify something *as* an anchor and then fail to make sufficient adjustments. The most difficult thing is realizing that something is serving as an anchor at all, much less that we are being influenced by it. We simply don’t recognize or represent the situation as one involving an anchor.

There are many other examples of our failure to identify situations as ones where specific cognitive tools could make our lives easier. For example, even people who avoid the gambler’s fallacy when thinking about the tosses of a coin may fail to ask whether a similar situation applies when we ask about the probability of Wilbur’s and Linda’s sixth child being a girl after an initial string of five boys (p. 311). And even if we have studied conditionals, we may have trouble with the card selection experiment (p. 551) because we don’t realize that it involves conditionals, necessary conditions, and the like.

Some of the questions we have encountered in the book, e.g., ones about the conjunction fallacy, are a little tricky, and they are often tricky *because* they make it difficult to *represent* the problem in the best way, e.g., as one about sample size or about the probability of a conjunction. But the fact that some of the problems seem tricky shouldn’t mislead us into thinking that they are just tricks in textbooks. Every day all of us fail to recognize or identify situations as ones in which various cognitive tool are applicable.

A good representation suggests the relevance of an appropriate cognitive tool

### 28.3.2 Cues That Signal when a Tool Applies

We could phrase the problem in various sound bites:

- Cognitive tools do not apply themselves.
- Situations do not come labeled as ones where a particular tool is applicable.

Knowing whether a particular concept or principle applies in a given situation can be very difficult, but there are a few *cues* that *signal* the relevance of a particular tool. Here are a few examples:

**Argument Indicators** One of the questions we face over and over in reasoning tasks is this: do these premises or reasons or assumptions support that conclusion? Such situations involve arguments, but we will never be able to apply the relevant tools for dealing with arguments (e.g., validity, conditional arguments, inductive strength, various fallacies) unless we recognize that a situation *involves* an argument.

Fortunately there are some obvious and easy cues that signal the presence of an argument. Conclusion indicators (e.g., ‘therefore’, ‘so’, ‘hence’) and premise indicators (e.g., ‘because’, ‘since’) tell us that we have probably encountered an argument, suggest which parts are premises and which conclusions, and give us a fighting chance of evaluating it.

**Necessary and Sufficient Conditions** Any time we encounter talk of requirements or prerequisites we are probably dealing with necessary conditions. And talk of guarantees or things that are enough usually signals sufficient conditions (p. 48).

**Probability Indicators** There are some obvious cues that signal the relevance of our concepts and rules involving probabilities. When words like ‘probable’ or ‘likely’ appear, or when we are dealing with cases everyone knows involve chance (the lottery, flipping a coin), it is usually clear that probabilities could be relevant.

Less obvious cues include words like ‘proportion’ and ‘percentage’ and ‘frequency’, since these can often be translated into probabilities. We also are dealing with probabilistic concepts and principles when we ask which of two (or more things) is more likely (is Wilbur more likely to have gone to the show or the bar?). And when we try to *decide* what to do or *predict* what will happen, we are often dealing with probabilities. Should I major in business or engineering? Well, if I major in business I’ll probably have an easier time finding a job, but on the other hand I probably won’t like the work quite as much.

Remember that we don't need precise numbers to use many of our probability tools. You reason better even if you only remember that—and why—disjunctions are typically more probable than their disjuncts and conjunctions are typically less probable (often *much* less probable) than their conjuncts.

**Why?** When someone (including you) makes a claim that isn't obviously true that can serve as a cue to ask *why?* What reasons are there to think that it's true? Are the good reasons to suspect that it's false? Getting in the habit of asking *why?*, when you encounter a claim, is one of the simplest things you can do to improve your reasoning. Not in an impolite way, where you quickly become a pain in the neck, or in the spirit of the child who will continue asking why no matter what answers we give. Ask in the spirit where you would accept a good answer.

**Realistic Tools** In the real-world we rarely have the time or energy or inclination to devote a huge amount of thought to the problems we encounter. We have better things to do on a Saturday night than sit home and calculate probabilities, but even if we enjoyed such things the limitations on human attention, working memory, and computational abilities mean that there are limits to what we can do. So *realistic* cognitive tools, one's we are likely to use when we need them, must be ones we can use without a great deal of time and effort.

We have encountered a number of ways to avoid specific fallacies and cognitive biases. In the next section we consider a strategy that is simple enough to be realistic, yet general enough to help us in a wide range of situations.

## 28.4 Consider Alternatives

We approach many situations with certain expectations or preconceptions, with a “mindset” that makes it easy to overlook the relevance of various cognitive tools and skills. We often think about a problem in the way it is presented to us, or in the way that we have gotten used to thinking about similar problems. Although there are no magic cures, there is a general strategy that can help overcome this obstacle to good reasoning. Consider Alternatives!

The basic idea is to think about a situation from several points of view or in several different ways, rather than to simply jumping to the first conclusion that occurs to us. Although this involves somewhat different things in different cases, it is easier to remember if we have a single, catchy label for the general strategy, so we will refer to it as “Consider Alternatives.”

In some cases consideration of alternatives means considering actual alternatives or options we have overlooked; in other cases it means considering hypothet-



ical alternative or options. All this sounds rather abstract, but we can bring it down to earth by considering a few concrete examples.

### Alternative Explanations

We constantly seek to explain the world around us. When something happens that matters to us we want to know *why*. What *explains* it? What *caused* it? Wilbur has been friendly to Wilma all semester but suddenly he begins insulting her. She naturally wonders why.

Unfortunately, we often jump to conclusions about causes and explanations. One corrective is to ask ourselves whether there might be alternative explanations that would explain what happened better than the explanation that first occurred to us.

**Regression Effects** We often overlook regression to the mean, the fact that more extreme outcomes or performances tend to be followed by ones closer to the mean (i.e., closer to average; p. 317).

We often do something, note what happens, and conclude that our action led to—*caused*—the outcome. For example, if the implementation of a new policy is followed by a decrease in something undesirable or an increase in something desirable, it is often tempting to conclude that the new measure caused the shift. Unemployment went up last year, this year the City Council lowered property taxes, and now unemployment has come back down to its normal level.

In at least some cases like this, however, the return to normal would have occurred *without* the new measure, simply as a result of regression to the mean. In such cases, we are likely to *explain* the drop in unemployment by the decrease in taxes, but we will be wrong and the new measure will be given credit it doesn't deserve. Here the alternative explanation to consider is that the return to a normal level is simply a result of regression to the mean; perhaps it would have happened anyway, even without the new policy.

**Illusory Correlation** Suppose that a number of people recover from a given illness after taking megadoses of Vitamin C. It is tempting to conclude that the vitamin led to—*caused*—their recovery, that it *explains* why they got better. But might they have gotten better anyway? In the case of many diseases and illnesses some people do get better without any medication or vitamins or treatment.

To know whether the Vitamin C helped, we need to compare the rate of recovery among those who took it with the rate among those who did not. In short, if we fail to consider the rate of recovery of those who don't take Vitamin C we



may come to believe in an illusory correlation between taking the drug and getting better.

In this case belief in an illusory correlation leads us to give a faulty explanation (the people got better *because* they took the vitamin). Here the *alternative explanation* we need to consider is that they would have gotten better *anyway*, without the vitamin, (and the *alternative evidence* we need to consider in order to test either explanation is the group of people who did not take the vitamin).

**Fundamental Attribution Error** The fundamental attribution error is the tendency to overestimate the role of peoples dispositions and traits, and to underestimate the importance of the context, in *explaining* their behavior (p. 479). For example, it is easy to believe that people in a crowd who don't help an injured person are selfish and uncaring, when in fact many kind and decent people are reluctant to help *in such situations*. There are various situational pressures—*causes*—that inhibit helping.

In cases like this we give a wrong, or at least one-sided, explanation when we explain what others do in terms of their character traits or other internal causes. Considering the alternative explanation, that there are strong situational pressures that would lead most people (including them) to do what they did, can help us see that an explanation weighted more in terms of external causes might be more accurate.

**Gambler's Fallacy** In the cases thus far we need to consider alternative explanations for events that have already occurred, but similar points can hold for events in the future. We commit the gambler's fallacy when we treat independent events as though they were dependent (p. 311). For example, even if we believe that a coin is fair, we may think that if we have flipped three heads in a row we are more likely to get a tail on the fourth flip.

Such a belief can result from accepting a bad explanation while ignoring alternative explanations about why we should *not* expect a tail the fourth time around. The bad explanation is that we think a tail is more likely because it would “even things out,” or satisfy the “law of averages.”

The good explanation we overlook is that there is nothing about the coin that would enable it to “remember” what it did on earlier flips. There is no mechanism that can change the probability of a given flip based on the outcome of previous flips. We might well see this quite easily if we paused and asked *how* the outcome of one flip *could* affect the outcome of the succeeding flip. What could *explain* how the coin could remember (answer: nothing—so it doesn't)?

### Alternative Possibilities

Sometimes asking how things might have turned out differently helps avoid fallacies and cognitive biases. Here we *consider alternative outcomes or possibilities*.

**Hindsight Bias** Hindsight bias is the tendency to overestimate the likelihood that one would have predicted an outcome after they learn that it occurred (p. 145). It has been found in judgments about elections, medical diagnoses, sporting events, and many other topics. It is also the key ingredient of Monday morning quarterbacking and second guessing.

As with many other biases and fallacies, simply warning people of the dangers of hindsight bias has little effect. But various studies suggest that we can reduce this bias by considering how past events *might have turned out differently*. Consider alternative scenarios: ask yourself what alternatives might have occurred and what things would have made it likely that they would have happened?

For example, if we consider how easily a couple of charging fouls against the opponent's star could have gone the other way, it may be easier to see how they could have won the game. Considering alternative outcomes can make it easier to break free of the "mindset" that the occurrence of this event would have been obvious before it actually occurred.

**Overconfidence** We have a tendency to overestimate the probability that our own judgments or predictions are correct. This means assigning high probabilities to hypotheses or claims that turn out, reasonably often, to be false. There is evidence that considering alternatives can help us make a more realistic assessment of our own accuracy.

For example, if we imagine plausible scenarios in which the candidate we predict to win would turn out to lose, we may become less confident in our prediction. Here the key is to consider ways ones views and beliefs might turn out to be wrong. This can help us break free of the belief that our belief or prediction is "obvious."

**The Either/Or Fallacy** We commit the either/or fallacy (the fallacy of a "false dilemma") when we assume that there are only two alternatives when in fact there are more or, more generally, when we assume that there are fewer alternatives than there actually are (p. 203). Here pausing to ask if we have overlooked genuine alternatives may help us find additional, and with luck better, options.

### Alternative Evidence

Sometimes we need to consider evidence (data, facts, information) that we have overlooked: here the strategy is to consider alternative (or, perhaps more accurately here, overlooked) evidence

**Testing** It often seems easy to explain how we were right, *whatever* the evidence turns out to be. We read earlier about the members of a doomsday cult who gave away all their possessions because they believed that the world would end on a given day (p. 379). Their belief turned out to be false—obviously so, since the world did not end when they said it would.

Most of us would see this as rock solid proof that they were simply mistaken about things. But when the appointed time came and went, many cult members actually strengthened their beliefs in the pronouncements of their leader (by concluding that their efforts had postponed the end). To take another example, the predictions of many pseudoscientists are often sufficiently fuzzy that they can be reconciled with almost *anything* that happens later.

If such cases it is useful to consider alternative ways that the evidence might have turned out. What outcomes, what possible evidence, would we have taken as disconfirming our views? How would we have reacted if the evidence *had* turned out differently? This can help us see whether we are basing our beliefs on evidence (which could turn out to disconfirm them), or whether we are simply holding them, “now matter what,” whatever the evidence might be.

### Alternative Frames

**Framing** Framing effects occur when the way in which options or situations are described influences how we think about them (p. 348). It is always a good policy to imagine alternative ways to frame things, especially frames in terms of gains (if the current frame is in terms of losses) or in terms of losses (if the current frame is in terms of gains). And when evaluating the arguments of others, it is always wise to ask yourself how the points they are making might be *reframed*.

**Anchors** We discussed anchoring effects earlier in this chapter. Here we simply note that considering alternative anchors can help us avoid insufficient adjustment to anchors that are skewed.

### Alternative Points of View

Sometimes the alternative we need to consider some other person's point of view. We may need to put ourselves in their shoes.

**Strawman** We commit the strawman fallacy if we distort someone else's position to make it easier to attack (p. 183). Here considering alternatives means asking how the person who defends the view we dislike would state her position. What reasons would *she* give to support it? The point here is just to play fair, to try to find the strongest version of the view in question and evaluate it. Consider the alternative, stronger, ways to defend it.

**Actor-Observer Asymmetry** We tend to see other people's behavior as internally caused, but to see our own as externally caused (p. 482). Wilbur donated money to the Cancer Society because he is kind and generous. By contrast, we tend to see ourselves as giving a donation because we think the cancer victims need help.

The actor-observer asymmetry is a bias in our reasoning about people's actions, both our own and those of others we observe. But it turns out that if an actor imagines herself in the situation of the observer, or if an observer imagines himself in the situation of the actor, this bias is reduced.

**Out-Group Homogeneity Bias** This is the tendency to see out-groups (e.g., people of other races, religions, countries, sexual-orientation) as more homogeneous than they really are (p. 500). Groups tend to see themselves as quite varied, whereas members of other groups are thought to be much more similar to one another.

Here one way to consider the alternative is simply to learn more about the other group. But imagining what the group must seem like to those in the group may also make it easier to see that they probably vary as much as the members of groups we are more familiar with. For example, people in familiar groups with quite different backgrounds, occupations, and ages often see things in different ways. So are members of an out group with quite different backgrounds, occupations, and ages really likely to see things in the same way?

**Dissent** We encountered several cases in Chapter 22—Ash's conformity studies, Milgram's studies of obedience—where the most effective way of reducing conformity or obedience was to have at least one person who refused to go along. Often even one dissenter was enough to eliminate groupthink or conformity or mindless compliance.

Here one way to consider alternatives is to imagine what a dissenter might say or do. It is also good to encourage the actual presentation of alternative points of view by encouraging free expression. Some group biases are less likely, for example, if group members are allowed to express disagreement, since the group will be exposed to alternative points of view. It will be easier to break the grip of a particular way of looking at the relevant issues.

**Evaluating Sources of Information** We have no choice but to rely on others for information, but some people are better sources than others. When someone who seems to be an expert makes a claim it is useful to consider perspectives or points of view from which the claims would seem less plausible, and ones from which it would seem more plausible. Then consider it from the source's own point of view: would he have any reasons to make this claim if he didn't think it was true?

### 28.4.1 Is there "Invisible Data"?

Good reasoning requires us to stay focused on things that are *relevant* to the points we are reasoning about.<sup>1</sup> Many pitfalls in reasoning, including some of those discussed earlier in this chapter, involve one of two kinds of mistakes about relevance. Sometimes we overlook or ignore evidence or facts that *are* relevant. And sometimes we rely on evidence or facts that are *irrelevant*. We might think of the first sort of error as one of omission (we overlook relevant information) and the second as one of commission (we incorrectly treat irrelevant information as though it were relevant).

We will begin with cases where we overlook or underutilize relevant data. The slogan here, the question to ask ourselves, is: "Is there invisible data?" Is there relevant data that is "invisible" to us because we overlooked it? For example, in the case of the falling cats, we may jump to the wrong conclusion because we fail to consider the cats that died (here the sample from the "cat morgue" is relevant; p. 274). Here are some more general examples of relevant information that we often simply fail to notice.

**Confirmation Bias** Confirmation bias is our tendency to look for, notice, and remember confirming or positive evidence (that supports what we think) while overlooking or downplaying disconfirming or negative evidence (which suggests that what we think is wrong; p. 359). We often have a blind spot for outcomes and possibilities at odds with our beliefs and expectations.

---

<sup>1</sup>Remember that information is not relevant all by itself, in isolation from anything else. Relevance involves a *relationship* between one statement and another. A piece of information can be highly relevant to one conclusion but completely irrelevant to others.

For example, if Wilma is convinced that people with red hair have bad tempers she may be more likely to notice or remember cases where red heads fly off the handle and to overlook or forget cases where they don't. Here the invisible data is obvious—*once we mention it*.

Here is another example. People who make admissions decisions at universities typically receive limited feedback on how well they are making their selections. They probably do learn about the performance of the students they do admit, but they rarely get feedback on the people they reject. How well did they do in the university they ended up attending? How well might they have done in this university, the one that rejected them. Here half of the relevant evidence is “invisible” to the admissions officers, so it is difficult for them to form an accurate assessment about how well they are doing.

In cases like this it is often difficult to get the invisible data, although there are a few exceptions like football recruiting, where someone who wasn't offered a scholarship goes on to be a star at another university. But often invisible data is easy to see, as with Wilma and red heads—once it occurs to us to look.

**Illusory Correlations** Belief in illusory correlations often results from considering only some of the relevant cases, e.g., the people who took Vitamin C and got better. We might learn that the apparent correlation is illusory if we also considered people who did not take Vitamin C (since at least as many of them might have gotten better too). The problem is that we often overlook or ignore data or information about this group, so it remains invisible *to us*. To take another example from an earlier chapter, if we looked for invisible data it might also show that various superstitions are based on illusory correlations.

**Insensitivity to Base Rates** We often ignore or underutilize information about base rates in judging probabilities and making predictions (p. 337). For example, we may think that a person has a high probability of being infected with the HIV virus after testing positive on an accurate, but not-perfect test, without taking into account the relatively low incidence of such infections in the general population (p. 554). Or we may think that someone has a certain job like being a professional football player because they fit the profile or stereotype we associate with that job, while ignoring the low base rate of professional football players in the general population (p. 335).

When we overlook base rates, we ignore relevant information; it remains invisible to us. It will help us make it visible if we pause to ask whether information about base rates is relevant. Remember that you usually do not need any precise knowledge of base rates. Just knowing that there are a lot more of one sort of thing

(e.g., bankers) than another (e.g., professional football players) is often enough.

**How Do we Know what might be Relevant?** It's easy for a book to tell us we shouldn't overlook relevant information, but how can we tell what all is relevant, especially in fields we don't know much about? Things are typically open-ended, and we could go on collecting evidence for years. We often have to act soon, however, and it's important to know when to stop looking and act.

There is no guideline for this, but in the real-world the problem is rarely that people spend too much time looking for relevant information. The problem is that we usually don't look enough. Indeed, often we don't look at all. Moreover, the information we need is often obvious, once we pause to think about things. In short, in many other cases fostering the habit of asking: "what's been omitted?" can help us reason more effectively.

### **Is some Data *Too* Visible?**

The flip side of overlooking relevant information is basing our reasoning on *irrelevant* information. Even worse, in some cases we *have* good information, but for one reason or another we don't use it. Indeed, bad information sometimes drives out good information. In cases like this the evidence here doesn't start out invisible, but bad or irrelevant information makes it difficult to keep it in sight (as with the dilution effect, p. 338).

We saw various cases on this sort in Chapter 10, where we examined various species of the fallacy of irrelevant reason. Later we saw that we often focus on samples that are highly available in memory or imagination, while ignoring more relevant and representative samples, even when we know about them (p. 331).

## **28.5 Acquiring Cognitive Skills**

The most important step toward better reasoning is learning to *spontaneously* recognize situations where a particular cognitive tool or skill should be applied. Simply learning to fill in a blank saying whether one sentence is a necessary or a sufficient condition for a second won't be of much help when you actually need to think about necessary and sufficient conditions in the real world. Situations (outside textbooks) do not come marked with labels saying that they involve necessary or sufficient conditions, so you need to be able to recognize them on your own. Similar points hold for most of the other concepts and tools we have learned about, including the very general strategies of considering alternatives or asking whether there is invisible data.



**Transfer of learning:** being able to apply knowledge and skills acquired in one setting in quite different settings

We encountered this general issue early in the book when we discussed inert knowledge (p. 145). Psychologists and educators also speak of **transfer of learning**—being able to transfer the concepts and principles you learn in one setting (like the classroom) to use them in other settings (on the job, in dealing with your family).

How *can* we teach—and learn—how to apply cognitive tools in life outside the classroom? These are empirical questions, and we are still discovering the answers. Experienced teachers know a good deal about such things, but they are not immune to the biases and mistakes discussed in earlier chapters. Fortunately, in addition to first-hand experience, there has now been a good deal of research on the transfer of cognitive skills, and much of it points to two conclusions:

- Teaching reasoning tools and skills requires *teaching for transfer*.
- Learning reasoning tools and skills requires *learning for transfer*.

### Teaching for Transfer

Teaching for transfer is absolutely critical

How do we teach for transfer? First, it is important that students *understand how* and *why* principles and rules work the way they do, rather than simply learning to apply them to simple examples by rote. Second students need a lot of practice applying the concepts and principles for critical reasoning in a wide range of situations, to a wide range of subject matters. If you want to be able to apply a skill (like noticing the relevance of base rates) in a wide range of circumstances, you need to practice it in a wide range of circumstances. If you want to be able to apply a skill (like spotting skewed anchors) in new and novel situations (like a new job), you need to practice it in new and novel situations.

It is especially important for students to apply the skills they are acquiring in situations that matter in their own lives. But it is unlikely that they will immediately see the relevance of some principles (e.g., regression to the mean) to such situations (e.g., ones involving choices of a major or job) without help.

It is also probably easier to gain a lasting mastery of cognitive tools if the course work (e.g., chapters exercises, exams) is cumulative. It is obvious that most of the exercises at the end of a section on necessary and sufficient conditions involve these concepts, and so the student is primed to use them. But real-life situations rarely come labeled to prime us in this way.<sup>2</sup>

It also appears that a variety of instructional tools are useful: lectures, individual projects, group discussion, group projects, computer projects, assembling a

<sup>2</sup>I haven't followed the advice about making exercises cumulative here, but I encourage assigning a few randomly selected problems from earlier chapters every so often



portfolio (to be graded) of actual situations the student encounters that involve the things learned in class.

### Habits

Much of what we do, we do out of habit—almost automatically, without giving it much thought. Good reasoning typically results from good cognitive habits (e.g., a habit of looking for relevant evidence before drawing a conclusion). Bad reasoning often results from bad ones (e.g., a habit of jumping to conclusions without considering the evidence).

This means that changing how we reason often requires changing habits, many of which are deeply ingrained, and learning for transfer requires enough practice that new and better habits can begin to take root. At first this requires conscious effort, but with further practice at least some reasoning skills can become more automatic, second-nature. This takes effort, and it requires real motivation to sustain it when there are more appealing things to do.

Although there is a good deal of research to support these points, the basic point is also a matter of common sense. No one supposes that you can learn to become a good cook just by reading cookbooks, without setting foot in the kitchen, or that you could learn to play the guitar merely by reading manuals. These are skills, and like all skills they can only be learned by practice. Lots of it. Critical reasoning is no different in this regard, but since it affects all aspects of our lives, it is more important.<sup>3</sup>

---

<sup>3</sup>For an entertaining, though sobering, discussion of several famous examples of self-defeating behavior see Barbara W. Tuchman's *The March of Folly: From Troy to Vietnam*. N.Y.: Ballantine, 1984. For a discussion of implications of the weaknesses in human reasoning for public policy see Chris Swoyer, "Judgment and Decision Making: Extrapolations and Applications," *Judgments, Decisions, and Public Policy*, Rajeev Gowda and Jeffrey C. Fox, eds., 2001, Cambridge University Press. What I have called "Consider Alternatives," is meant as a generalization of the strategy of "considering the opposite" advocated by Charles G. Lord, Mark R. Lepper and Elizabeth Preston, "Considering the Opposite: A Corrective Strategy for Social Judgment. *Journal of Personality & Social Psychology*, 47(6), 1984; 1231-1243. Craig A. Anderson and Elizabeth S. Sechler, "Effects of Explanation and Counterexplanation on the Development and Use of Social Theories," *Journal of Personality & Social Psychology*, 50(1), 1986; 24-34 provides additional evidence that this strategy is often effective. For discussions of transfer of learning and the teaching of reasoning skills see John R. Anderson, L. M. Reder, and Herbert A. Simon, "Situated Learning and Education," *Educational Researcher*, 1996, 25; 5-11, and Richard E. Nisbett, *Rules for reasoning*, 1993, Hillsdale, NJ: Erlbaum Associates. Further references to be supplied.

## 28.6 Chapter Exercises

1. The bad news is that it will often be easier to apply lessons about the obstacles to critical thinking to exploit human fallibility than to reduce it. For example, when policy making is poll driven, shifting the views of two or three percent of the electorate may determine what law gets enacted or who gets elected. A thirty-second TV spot warning people about regression effects is unlikely to have much impact on this, but an attack add that frames an issue in a frightening way very well may (of course this is all the more reason to learn to think well—as self-protection).

Give an actual example (from your own life, TV, some other course) where someone has done this. Explain the situation in detail, then analyze the tendencies to faulty reasoning that it exploits.

2. Earlier in this chapter we listed three examples (Clinton, Enron, the Bay-of-Pigs) where bad reasoning and decision making led to serious troubles. List another example, either from the news or from history. Explain the ways that some of the pitfalls we have studied may have been involved in the case you describe.

3. Earlier in this chapter we noted three pitfalls in reasoning that can make it difficult to get accurate feedback about the accuracy of our reasoning: hindsight bias, confirmation bias, and wanting to avoid unpleasant feedback. Give another way one of the biases or fallacies we studied could make it difficult for us to get accurate feedback and explain why.

4. You are a teacher, and one of the topic you must cover is critical reasoning. What would you do to combat the problem of inert knowledge? Pick the level of the class (e.g., sixth grade, junior in high school, college freshman), but say what the level is.

A. Give a specific, detailed (half a typed page) example of an assignment on the concept of illusory correlation that is designed to give your students a working (rather than merely inert) grasp of the concept. Don't just say that you would ask students to give real-life examples. Give one or two examples of the sort of that they might encounter and find interesting, then say what additional directions you would give.

B. Give a specific, detailed (half a typed page) example of an assignment on the fundamental attribution error that is designed to give your students a working (rather than merely inert) grasp of the concept. Don't just say that you would ask students to give real-life examples. Give one or two examples of the sort of that they might encounter and find interesting, then say what additional directions you would give.

5. One of the problems we encountered in earlier chapters is the card-selection problem.

Wilbur has a pack of cards, each of which has a letter [either a consonant or a vowel] on one side and a number [either even or odd] on the other. Some of the cards lying flat on a table (Figure 27.3 on page 551). Which cards should you turn over in order to help test the hypothesis:

Hypothesis: If a card has a vowel on one side, then it has an odd number on the other.

As we saw (p. 551), the answer is that only two sorts of cards are relevant: vowels (most people get this right) and even numbers (many people get this wrong). This is a problem involving the basic features of conditionals, but even after studying them many people find the problem difficult. Why do you think this is so? What things could you do (over and above those mentioned in class) to help someone see what the correct answer is?

6. Good performances are likely to be followed by less outstanding performance simply because of regression, and unusually bad performances by better ones. Discuss the ways in which this could lead us to make flawed predictions and to give faulty explanations.

7. Recall this exercise from an earlier chapter: Suppose that you are a good chess player and that Wilbur is good, but not as good as you. Would you be more likely to beat him in a best of three series or in a best of seven series (or would the number of games make any difference)? Like the hospital problem above, this is a problem about sample size, but many of us fail to realize this when we first encounter the problem. First explain what the correct answer is and say why it is correct. Then discuss reasons why it might be difficult to identify this problem as one that involves sampling.

### Answers to Selected Exercises

6. Hint: if we neglect the possibility of regression effects, this may lead us to suppose that criticizing someone for a bad performance is a more effective way to getting them to do well than praising them for a good performance is. How could this encourage us to make flawed predictions and to give faulty explanations?



# Appendix: Pretest

Circle the number that you think gives the best answer to each question. Don't agonize over the questions, but give each a serious answer. Disregard the numbers in braces (e.g., [8]); this just refers to the chapter that deals with the question.

1. Suppose that you toss a fair coin a number of times. Which of the following sequences of heads (H) and tails (T) is most likely? [13 & 16]

1. HHHHHTTTTT
2. HTTHTHHHTH
3. They are equally likely

2. Suppose someone had fallen down and was bleeding. Which of the following do you think is most likely? [23]

1. People in general would be most likely to help if they were the only person around.
2. People in general would be most likely to help if they were several other people around.

3. Are there more 6 letter English words (a) ending in 'ing' or (b) with 'n' as their fifth letter? [16]

1. More ending in 'ing'
2. More with 'n' as their fifth letter
3. There are the same number of each sort of word
4. It is impossible to tell without counting all of the words in English.

4. If you were a juror you should [7 & 8]

1. place a good deal of weight on eyewitness testimony

2. place some, but not a huge amount, of weight on eyewitness testimony
  3. not place much weight on eyewitness testimony at all
5. Which alternative seems most likely to occur in the next ten years? [16]
1. An all-out nuclear war?
  2. An all-out nuclear war that accidentally develops out a confrontation in the Middle East involving Iraq or Iran and some of their neighbors and that then spreads out of the region to other countries?
6. Which is the more likely cause of death? [17 & 21]
1. homicide
  2. suicide
7. Which is the more likely cause of death? [17 & 21]
1. fire
  2. downing
8. About how many people would need to be sampled in order to make a reasonably accurate prediction (say within 5 percentage points) about who U.S. voters prefer for President (suppose that there are about 100 million voters)? [15]
1. about five thousand
  2. about five hundred thousand
  3. about five million
  4. more than five million
9. Suppose an normal student signed up to be a subject in a psychological experiment. When he got there, the experimenter asked him to administer a series of ever more painful shocks to another subject who showed up at the same time. [22]
1. Most subjects would not administer any shocks at all
  2. Most subjects would just administer a few shocks but quit before they got to the really painful ones
  3. Most subjects would administer shocks but quit as soon as the other subject screamed in pain.
  4. Many subjects would continue to administer ever more painful shocks, even after the other subject screamed in pain and pleaded he had a heart condition.

10. Sally and Bob have had five children, all of them boys. They would like very much to have a girl. [16]

1. They should go ahead and have another child, because by the “law of averages” things are likely to even out, and they will probably have a girl.
2. They have always had boys, so they would probably have a boy again
3. The chances of having a girl and the chances of having a boy are virtually the same.

11. Which of the following alternatives—1 or 2—do you prefer? [18]

1. a 100% chance of losing \$50
2. a 25% chance of losing \$200, and a 75% chance of losing nothing.

12. There are two hospitals in Belleville. About 50 babies are born everyday in the larger one, and about 15 are born everyday in the smaller one down the street. On average 50% of the babies are boys, but the number bounces around some from day to day. Which hospital, if either, is more likely to have more days per year when over 65% of the babies born are boys? [15]

1. The big hospital
2. The small hospital
3. The two hospitals should be about the same

13. Suppose that you have a pack of special cards, each of which has a letter [either a consonant or a vowel] on one side and a number [either even or odd] on the other. If you have some of the cards lying flat on a table, which ones would it be relevant to turn over if you were trying to *disprove* the following claim? [3]

Cards with vowels on one side always have odd numbers on the other side.

1. cards with consonants and cards with even numbers on them.
2. cards with vowels and cards with even numbers on them.
3. cards with consonants and cards with odd numbers on them.
4. cards with vowels and cards with odd numbers on them.
5. you need to turn over *all* of the cards.
6. None of the above.

14. Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations. Which of the following is most likely? [16]

1. Linda is a bank teller.
2. Linda is a bank teller and is active in the feminist movement.

15. Exactly two cab companies operate in Belleville. The Blue Company has blue cabs, and the Green Company has Green Cabs. Exactly 85% of the cabs are blue and the other 15% are green. A cab was involved in a hit-and-run accident at night. A witness, Wilbur, identified the cab as a Green cab. Careful tests were done to ascertain peoples' ability to distinguish between blue and green cabs at night. The tests showed that people were able to identify the color correctly 80% of the time, but they were wrong 20% of the time. What is the probability that the cab involved in the accident was indeed a green cab, as Wilbur says? [17 & 28]

1. Very likely, almost 100%
2. The chances are about 85%
3. The chances are about 40%
4. Less than 10%

16. Which conclusion (a, b, etc.) follows best follows from premises 1 and 2: [2]

1. No members of the ad-hoc committee are women.
2. Some U.S. Senators are members of the ad-hoc committee.  
Therefore
  - (a) All U.S. Senators are women.
  - (b) No U.S. Senators are women.
  - (c) Some U.S. Senators are women.
  - (d) Some U.S. Senators are not women.
  - (e) None of the above really follow from the two premises.

17. The grade school in Belleville administers an achievement test to all the children who enter fifth grade. At the end of the school year they give the same test again. The average score each time is 100, but something odd seems to have happened. Children who scored below average on the test the first time tend to improve (by about five points), and children who scored above average tend to do worse (by about five points). What's going on? [16]

1. When the two groups interact, their members influence the other group (high scorers pull up low scorers and low scorers pull down high scorers).
2. The two groups probably don't affect each others abilities very much, but each group may affect they way the other group takes tests.



3. This is simply the sort of thing one should expect in a situation like this.

18. Mike is six two, muscular, and in excellent physical condition. He was a college athlete and is aggressive. Although one can't be sure, which of the following two is more likely. [17]

1. Mike is a banker.
2. Mike is a pro football player.
3. They are about equally likely.

19. Visual perception [4]

1. works very much like a video camera (or video tape recorder)
2. is not much like a video camera at all

20. How likely (probable) is it that your vote could determine the outcome of a U. S. Presidential election? [13 & 26]

1. Fairly likely
2. Somewhat likely
3. Not very likely, but still possible
4. No real chance whatsoever

21. There are three doors in front of you. There is nothing worth having behind two of them, but there is a suitcase containing \$50,000 behind the third. If you pick the correct door, the money is yours. You choose door number 1. But before the host of the game show shows you what is behind that door, he opens one of the other two doors, picking one he knows has nothing behind it. Suppose he opens door number 2. This takes 2 out of the running, so the only question now is about door 1 and door 3. He allows you to reconsider your earlier choice: you can either stick with door 1 or switch to door 3. What is the probability that the money is behind door 3? [14 & 27]

1.  $1/3$
2.  $2/3$
3.  $1/2$
4. There is no way to tell.

22. Which conclusion (a, b, etc.) follows best follows from premises 1 and 2: [2]

1. No U.S. governors are members of the Harem Club
2. Some Arabian shieks are members of the Harem Club

Therefore

- (a) All Arabian shieks are U.S. governors.
- (b) No Arabian shieks are U.S. governors.
- (c) Some Arabian shieks are U.S. governors.
- (d) Some Arabian shieks are not U.S. governors.
- (e) None of the above really follow from the two premises.

23. Wilbur is on his way to class and encounters someone slumped over in a doorway. The person is obviously in distress and needs help. Suppose you had to predict whether Wilbur would stop and help. Which single piece of information would help you make the most accurate prediction? [23]

1. You would need to know a good deal about Wilbur's personality and traits to make a very good prediction
2. You would need to know if Wilbur was in a hurry or not.

24. Suppose that you are an excellent chess (or squash) player and Wilbur is good, but not as good as you. [16]

1. You would have the best chance of beating Wilbur in a best of seven series
2. You would have the best chance of beating Wilbur in a best of three series
3. Your chances of beating Wilbur would be the same in either case.

25. It takes about eleven days for a million seconds to tick off. Roughly how long does it take for a billion seconds to tick off? [16 & 20]

1. About 8 to 10 weeks
2. Between 6 and 11 months
3. 2.5 years
4. 32 years
5. 8000 years

26. A number of college students participated in an experiment. They were divided into two groups. Each subject was tested with just themselves and an experimenter present. The students in the first group were asked nicely by a warm, friendly experimenter to eat several fried grasshoppers on a plate in front of them (in order to help her gather data for her dissertation). Those in the second group were pressured by a cold, aloof experimenter eat several grasshoppers. Most subjects in both

groups ate several grasshoppers. Later they were asked by a third person how much they liked the grasshoppers. What do you think happened? [19]

1. Those who had been asked by the nice friendly experimenter said that they liked the grasshoppers more than those in the other group.
2. Those who had been pressured by the cold, unfriendly experimenter said that they liked the grasshoppers more than those in the other group.
3. On average, subjects in the two groups reported liking the grasshoppers about the same.

27. A group of men (the members of a club in Belleville) consists of 70 engineers and 30 lawyers. Suppose that we select Dick at random from the group. The following is true of Dick:

Dick is a 30 year old man, married, with no children. He has high ability and high motivation, and promises to be quite successful in his field. He is well liked by his colleagues.

Based on this: [17]

1. It is more likely that Dick is an engineer.
2. It is more likely that Dick is a lawyer.
3. Dick is equally likely to be an engineer or a lawyer.

28. When it comes to learning the material in a course [7 & 8]

1. Tape recording the lectures is a good way to learn.
2. Tape recording the lectures is a hopeless way to learn.
3. Tape recording the lectures works about as well as most other methods.

29. Which of the alternatives, 1 or 2, do you prefer? [18]

1. an insurance policy with a \$50 premium that protects you against losing \$200
2. a 25% chance of losing \$200, and a 75% chance of losing nothing.

30. In five seconds or less give you best estimate of:

$1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$  [17]

1. about 512
2. about 2250
3. about 11,00
4. about 40,000
5. about 200,000.



# Index of Key Concepts

- actor-observer asymmetry, 482
- ad hominem fallacy, 180
- affirming the antecedent, 53
- affirming the consequent, 54
- anchoring and adjustment, 340
- appeal to ignorance, 187
- appeals to emotion
  - illegitimate, 161
  - legitimate, 160
- argument against the person, 180
- arguments, 24
  - conclusions of, 24
  - in standard form, 27
  - premises of, 24
- assurance problem, 527
- availability heuristic, 331
  
- bandwagon effect, 97
- base rates, 337
- Bayes' rule, 571
- begging the question, fallacy of, 198
- belief perseveration, 144
- bias, two meanings of, 506
- blind studies, 401
- bookmarks, in a browser, 109
- bottom-up processing, 68
- bright lines, 353
- broken-leg cases, 412
- burden of proof, 188
  
- central tendency
  - measures of, 275
  
- certainty effect, 352
- cognitive tools, 17
- collective memory, 148
- common cause, 292
- compromise effect, 367
- conclusion indicators, 26
- conditional, 43
  - antecedent of, 44
  - argument, 53
  - consequent of, 44
- conditional probability
  - definition of, 251
- confirmation bias, 359
- confounding of variable, 400
- conjunction fallacy, 312
- conservative updating of probabilities, 578
- constructive dilemma, 209
- contrast effect, 366
- control group, 400
- correlation, 287
  - and probability, 288
  - diagrams for, 290
  - illusory, 295
  - vs. causation, 292
- counterexamples, 34
- cues
  - to applicability of tool, 592
  
- De Morgan's Laws, 568
- decision trees, 566
- deductive validity, 30

- denying the consequent, 54
- denying the antecedent, 55
- diagrams, 545
- dilution effect, 338
- disjunctions, 203
  - expressed as conditionals, 207
- disjunctive syllogism, 208
- dispersal, measures of, 276
- double blind studies, 401
- doubling time, 424
  
- effect sizes in experiments, 364
- either/or fallacy, 203
- empirical questions, 7
- endowment effect, 351
- ethnocentrism, 501
- expected value, 308
- experiment, 400
  - experimental group, 400
  - experimental manipulation, 401
- expertise, 88
  - evaluating claims to, 89
- explanation reflex, 10
  
- fair bets, 259
- fallacies, 178
  - summary of, 217
- false dilemma, 203
- false memories, 142
  - false-memory syndrome, 143
- framing
  - as gains, 349
  - as losses, 349
- framing effects, 348
- free riders, problem of, 526
- freedom of expression, 466
- fundamental attribution error, 479
  
- gambler's fallacy, 311
- group polarization, 488
- groupthink, 489
  
- halo effect, 94
- hasty generalization, 283
- Hawthorne effect, 409
- heuristics, inferential, 330
- hindsight bias, 145
- hostile media phenomenon, 80
- hypertext, 108
  
- illusory correlation, 295
- incompatibility, 236
- inconsistency, 213
- independence, 240
- indicator words, 26
- inductive strength, 39
- inert knowledge, 145
- inference, 24
- information costs, 90
- information processing, 68
- intuitive vs. statistical prediction, 409
- inverse probabilities, 339
- irrelevant reason, fallacy of, 178
  
- jig-saw classroom, 516
- just-word hypothesis, 364
  
- Lake Wobegon effect, 164
- loss aversion, 351
  
- magic numbers, 353
- mean, 275
- media effects, 333
- median, 275
- memory
  - stages of, 126
- mere exposure, 362
- Mill's methods, 396
- misattribution of source, 140
- misinformation effect, 141
- mode, 275
- multiple sources
  - importance of, 100

- necessary condition, 48
- neglected evidence, 191
- null hypothesis, 402
  
- odds, 259
- out-group homogeneity bias, 500
  
- Pascal's wager, 310
- paternalism, 451
- perceptual inference, 70
- perceptual set, 73
- persuasion
  - by argument, 450
  - by manipulation, 450
- population
  - often difficult to determine, 401
- populations, 274
- post hoc, ergo propter hoc, 394
- preference reversals, 350
- premise indicators, 27
- primacy effect, 147
- prisoner's dilemma, 522
  - iterated, 533
- probabilistic reasoning
  - without precise numbers, 307
- probability
  - conditional, 249
  - inverse probabilities, 339
  - of conjunctions
    - with dependent conjuncts, 253
    - with independent conjuncts, 241
  - of disjunctions
    - with compatible disjuncts, 243
    - with incompatible disjuncts, 236
  - of negations, 234
  - rules for calculating, 230
  - specificity lowers, 314
  - summary of rules, 267
- professional persuaders, 454
- pseudo science, 414
  
- psychological accounting, 352
- public goods, 525
- Pygmalion effect, 361
  
- racism
  - modern, 498
  - overt, 498
- random assignment of subjects, 400
- random sample, 281
- recency effect, 147
- reference class
  - problem of determining, 576
- regression to the mean, 317
- relevance, 174
- representativeness heuristic, 335
- reversing figures, 70
- risk, 428
  - misperception of, 437
  - risk ratios, 429
- risk aversion, 349
- risk seeking, 349
- rule of 70, 424
  
- samples, 274
  - biased, 280
  - good, 279
  - random, 281
  - stratified random, 281
- sampling, 278
- sampling distribution, 402
- sampling variability, 282
- schemas, 135
- search engine, 110
- self-fulfilling prophecy, 361
- self-perception theory, 376
- self-serving biases, 163
- set, perceptual, 73
- significance testing, 402
- skewed searches, 118
- social influences on thinking, 452

soundness, of arguments, 33  
status quo bias, 351  
stealth advocacy, 119  
stereotypes, 495  
strawman fallacy, 183  
stress  
    effects on reasoning, 159  
subtyping, 505  
sufficient condition, 48  
sunk costs, 358  
suppressed evidence, 191  
  
top-down processing, 75  
total probabilities, rule for, 573  
transfer of learning, 601  
trees, probability, 557  
  
ultimate attribution error, 515  
  
validity effect, 362  
validity, of arguments, 30  
Venn diagrams  
    muddy, 235  
visual thinking, 545  
Voodoo science, 414