



Australian Government

BETA

BEHAVIOURAL ECONOMICS TEAM OF
THE AUSTRALIAN GOVERNMENT

Going blind to see more clearly: unconscious bias in Australian Public Service shortlisting processes

June, 2017

Authors: Professor Michael J. Hiscox, Tara Oliver, Michael Ridgway,
Lilia Arcos-Holzinger, Alastair Warren and Andrea Willis



© Commonwealth of Australia 2017

978-1-925362-52-7 Going blind to see more clearly (DOCX)

978-1-925362-51-0 Going blind to see more clearly (PDF)

Copyright Notice

With the exception of the Commonwealth Coat of Arms, this work is licensed under a Creative Commons Attribution 4.0 International licence (CC BY 4.0)

(<http://creativecommons.org/licenses/by/4.0/deed.en>).



Third party copyright

Wherever a third party holds copyright in this material, the copyright remains with that party. Their permission may be required to use the material. Please contact them directly.

Attribution

This publication should be attributed as follows: Commonwealth of Australia, Department of the Prime Minister and Cabinet, *Going blind to see more clearly*.

Use of the Coat of Arms

The terms under which the Coat of Arms can be used are detailed on the following website:

<http://www.itsanhonour.gov.au/coat-arms/>.

BETA

BEHAVIOURAL ECONOMICS TEAM
OF THE AUSTRALIAN GOVERNMENT

Going blind to see more clearly: unconscious bias in Australian Public Service shortlisting processes

Results of a randomised controlled trial conducted by the Behavioural Economics Team of the Australian Government (BETA) in partnership with the Australian Public Service Commission

Other uses

Enquiries regarding this licence and any other use of this document are welcome at:

Managing Director

Behavioural Economics Team of Australia

Department of the Prime Minister and Cabinet

Barton ACT 2600

Email: beta@pmc.gov.au

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Department of the Prime Minister and Cabinet or the Australian Government.

Acknowledgments

Special thanks to Professor Iris Bohnet (Harvard Kennedy School) and Kate Glazebrook (The Behavioural Insights Team) for feedback and suggestions. Our sincere thanks go to all of the people who generously donated their time to participate in the trial, to members of the working group from each participating agency who helped us understand their agencies better, and without whom this trial would not have been possible.

The trial was pre-registered on the BETA website and the American Economic Association registry: AEARCTR-001783. <https://www.socialscienceregistry.org/trials/1783/history/11939>

WHO WE ARE

We are the Behavioural Economics Team of the Australian Government, or BETA.

We are the Australian Government's first central unit applying behavioural economics to improve public policy, programs and processes. Rather than expecting people to redesign their lives around government, our work encourages people-centred design, which means: simpler, clearer and faster public services.

We use behavioural economics, science and psychology to improve policy outcomes.

Our mission is to build behavioural economics capability across the public service and drive its use in policy design by testing what works, where and in what context.

Contents

| | |
|---|----|
| Foreword | 5 |
| Summary | 6 |
| Why was it important to conduct this trial? | 7 |
| What interventions were tested? | 9 |
| How was the trial designed? | 11 |
| Results of the trial | 13 |
| Limitations of the trial | 16 |
| Policy lessons | 17 |
| Endnotes | 18 |

FOREWORD

We are excited to share the findings from our trial testing the impact of de-identifying applications for senior positions in the Australian Public Service (APS). The trial examines one critical part of the recruitment process, the way the APS initially evaluates applications and shortlists the top candidates to build the best possible workforce. The aim was to assess whether de-identifying applications would, by eliminating the effects of explicit or implicit bias, help promote gender equality and diversity in hiring at senior (executive) levels.

What we found is that de-identifying applications at the shortlisting stage of recruitment does not appear to assist in promoting diversity in hiring. In fact, in the trial we found that overall, APS officers generally discriminated in favour of female and minority candidates. This suggests that the APS has been successful to some degree in efforts to promote awareness and support for diversity among senior staff. It also means that introducing de-identification of applications in such a context may have the unintended consequence of decreasing the number of female and minority candidates shortlisted for senior APS positions, setting back efforts to promote more diversity at the senior management levels in the public service.

This is critically useful knowledge. It does not imply that the APS has solved the problem of gender equality at the executive levels and higher – or lack of diversity more generally – but it tells us that rather than putting the focus on bias in initial reviews of job applicants, it may be more valuable to direct attention to other stages of recruitment, including how positions are advertised, how interviews are conducted, and how hiring panels are selected and run. More attention can also be directed to processes that affect career trajectories, including performance reviews, evaluations for promotions, talent management and whether flexible working arrangements are available. Eliminating or mitigating problems in these areas will require innovative solutions and rigorous testing to discover what works.

Professor Michael J. Hiscox
Director, BETA

SUMMARY

This study assessed whether women and minorities are discriminated against in the early stages of the recruitment process for senior positions in the APS, while also testing the impact of implementing a 'blind' or de-identified approach to reviewing candidates.

Over 2,100 public servants from 14 agencies participated in the trial¹. They completed an exercise in which they shortlisted applicants for a hypothetical senior role in their agency. Participants were randomly assigned to receive application materials for candidates in standard form or in de-identified form (with information about candidate gender, race and ethnicity removed).

We found that the public servants engaged in positive (not negative) discrimination towards female and minority candidates:

- Participants were 2.9% *more* likely to shortlist female candidates and 3.2% *less* likely to shortlist male applicants when they were identifiable, compared with when they were de-identified.
- Minority males were 5.8% *more* likely to be shortlisted and minority females were 8.6% *more* likely to be shortlisted when identifiable compared to when applications were de-identified.
- The positive discrimination was strongest for Indigenous female candidates who were 22.2% more likely to be shortlisted when identifiable compared to when the applications were de-identified.

Interestingly, male reviewers displayed markedly more positive discrimination in favour of minority candidates than did female counterparts, and reviewers aged 40+ displayed much stronger affirmative action in favour for both women and minorities than did younger ones.

Overall, the results indicate the need for caution when moving towards 'blind' recruitment processes in the Australian Public Service, as de-identification may frustrate efforts aimed at promoting diversity².

WHY WAS IT IMPORTANT TO CONDUCT THIS TRIAL?

Women are almost 50% of the Australian workforce, but are under-represented in management and executive level positions. This is evident not just in the private sector but also in many areas of the APS. In 2016, women comprised 59.0% of the APS as a whole, but accounted for 48.9% of its executive level officers and only 42.9% of its Senior Executive Service (SES) officers.

These statistics may reflect a range of factors, including the availability and adoption of flexible work arrangements that make it easier for individuals to balance the demands of work and family. The numbers may also reflect gender discrimination in hiring and promotion. Such discrimination can be especially difficult to overcome when it results from unconscious cognitive biases (see Box 1) that are so internalised that people are unaware that their decision-making processes are affected³.

BOX 1: COMMON BIASES AFFECTING RECRUITMENT

Affinity / Ingroup bias may lead recruiters to prioritise candidates who are similar to them or someone they know and enjoy working with

Confirmation bias is the tendency to focus on information that confirms initial impressions of a candidate

Groupthink can occur when members of a recruitment panel feel pressure to conform with the decision of other panel members and may suppress their own opinions so as not to disturb the perceived group consensus

Halo effect may lead recruiters to focus on salient pieces of information and this may influence the perception of other elements of a candidate's application

Status quo bias can occur when recruiters opt for the 'safer' choice of recruiting a candidate that is similar to previously hired candidates, then the riskier option of hiring a candidate with different characteristics and/or background

In 2016 the APS announced a comprehensive [Gender Equality Strategy](#) aimed at addressing gender imbalance across all agencies and supporting a range of initiatives, including training and awareness programs to help promote diversity. In this context, considerable attention has been focussed on de-identification of job applications (and 'gender-blind' processes in particular) as an approach that could help reduce discrimination and promote diversity in the APS. A number of APS agencies have experimented with 'blind' job application processes. However these changes have not been consistently implemented, nor the impact rigorously measured.

Unfortunately, the existing evidence on the effects of de-identifying job applications is limited and mixed. Several quasi-experimental studies in European countries have suggested de-identification could reduce bias in hiring

processes in some contexts, but may have no impact in contexts in which no discrimination is present initially and, more perversely, may actually undermine efforts to promote diversity when employers adopt a positive bias in favour of women or minorities. To date there have been no rigorous studies designed to assess the extent and nature of bias affecting recruitment into the APS or to evaluate the impact of introducing de-identification processes. A summary of one of the first studies on the impact of ‘blind’ processes is provided in Box 2.

BOX 2: CASE STUDY – THE IMPACT OF ‘BLIND’ AUDITIONS ON FEMALE MUSICIANS

In the 1970s and 1980s American symphony orchestras attempted to overcome biases in hiring by introducing a screen during auditions to conceal the identity of the musician from the jury evaluating the performance. In a well-known study analysing data on auditions and hiring by orchestras over this period, this study found that the use of blind auditions had a major impact on gender bias in orchestras, increasing the likelihood of female musicians being selected by 25-40%⁴.



WHAT INTERVENTIONS WERE TESTED?

Proponents of de-identifying job applications argue that even when members of hiring and promotion committees are trained to be attentive to gender equity and potential discrimination, implicit or unconscious bias can still play a large role and weigh against female and minority job candidates⁵. Recruiters and reviewers are influenced implicitly by stereotypes when making judgements and often favour candidates with similar characteristics to their own. The one sure way to eliminate these types of bias is to disable the 'fast-thinking' processes in the minds of reviewers who rely on these types of heuristics by removing information about the characteristics of the candidates that are not relevant to potential performance on the job (see Box 3 for the definition of a heuristic). Rather than focusing on retraining the mind, we have redesigned the process in an attempt to eliminate unconscious bias.

BOX 3: WHAT IS A HEURISTIC?

Heuristics are mental shortcuts we use to ease the cognitive load of making a decision. Taking shortcuts by using a 'rule of thumb' (e.g., choosing a 50/50 mix of two options, or "one of each", or the "middle of the road" option), making an 'educated guess' based on past experience, or even just looking at what other people are doing are all examples of applications of heuristics.

There are many potentially irrelevant characteristics that could be screened out from reviewers in order to remove biases. Besides gender, race or ethnic status, we might also consider any information about age, health status or conditions, disability, sexual orientation, political views, and socioeconomic status (reflected for instance, by address and education background). De-identification could be implemented at the initial stage of the review process when applications are shortlisted and scored, but could also be extended to later stages (e.g., recruitment committee deliberations) – although it is obviously quite difficult (albeit not impossible) to de-identify candidates at the interview stage. In the current study we focussed on removing information about the gender, race and ethnic status of job candidates from written materials available to reviewers during the initial shortlisting stage of the APS recruitment process.

We developed a set of applications for an executive level APS position, with each application consisting of a 2-page curriculum vitae (CV) for each candidate. We standardised the format in which information in each CV was presented, allowing for no variation in font size, type, or colour and no other graphical features (e.g., photographs), and placing all content in the same order. We created a de-identified version of each application, in which the candidate was referred to only by number, along with male and female versions in which the candidate was referred to by the same last name and either a male or female first name. Several of the names were associated with specific minorities (Indian, Chinese, and Middle Eastern) and one candidate was explicitly identified as being of Indigenous descent (via a checked box at the top of the CV)⁶.

By examining the way reviewers evaluated the same applications with no identifying information compared with when they had information about the gender and race or ethnic status of the candidates we are able to assess the extent and direction of bias as well as the impacts of introducing de-identification at this stage of the recruitment process. For example, we looked at how CV1 was ranked and then compared this to the same CV when it was presented with a male/female name to determine any differences in how it was assessed and ranked.



HOW WAS THE TRIAL DESIGNED?

The trial was an individually randomised controlled trial (RCT) – more precisely, a ‘framed field experiment’ – in which subjects were drawn from the population of interest and invited to take part in an exercise in a natural (‘field’) setting for research purposes. In this case, a sample of executive level (EL) and Senior Executive Service (SES) Band 1 officers within 14 APS agencies were invited to participate in the study. Participants were asked to assess 16 hypothetical candidates for an EL position on a hypothetical taskforce and select a shortlist of 5 candidates. The CVs described a set of 16 realistic candidates with varied characteristics in terms of education and work experience such that shortlisting task was challenging for reviewers. An online survey platform was used to complete the exercise.

BOX 4: WHAT IS A RANDOMISED CONTROLLED TRIAL (RCT)?

A randomised controlled trial (RCT) is the best way of telling if a policy is working. RCTs work by randomly assigning individuals (or other units) into different groups – usually one or more ‘treatment’ groups that participate in the new intervention, and a ‘control’ group that does not. The differences in outcomes across the groups are then compared. RCTs are considered the ‘gold standard’ for assessing causal impacts because a RCT determines the impact of an intervention or treatment compared to if nothing was changed.

More information about RCTs is available [here](#).

Participants were randomly assigned to one of three groups to review the CVs in either de-identified (treatment) or usual identified (control) format. There were 2 control groups, each with 8 candidates identified as women and 8 as men; the only difference between the 2 control groups was that the first names used for the CVs in control group 1, were substituted with a similar first name of the opposite gender in control group 2 (e.g. the name Gary Richards in control group 1 became Wendy Richards in control group 2). To test for minority bias, in each control group there were 3 minority-sounding names included and 1 candidate was identified on their CV’s as being Indigenous.

FIGURE 1: TRIAL DESIGN - STANDARDISED CV CONTENT FOR CONTROL AND TREATMENT GROUPS



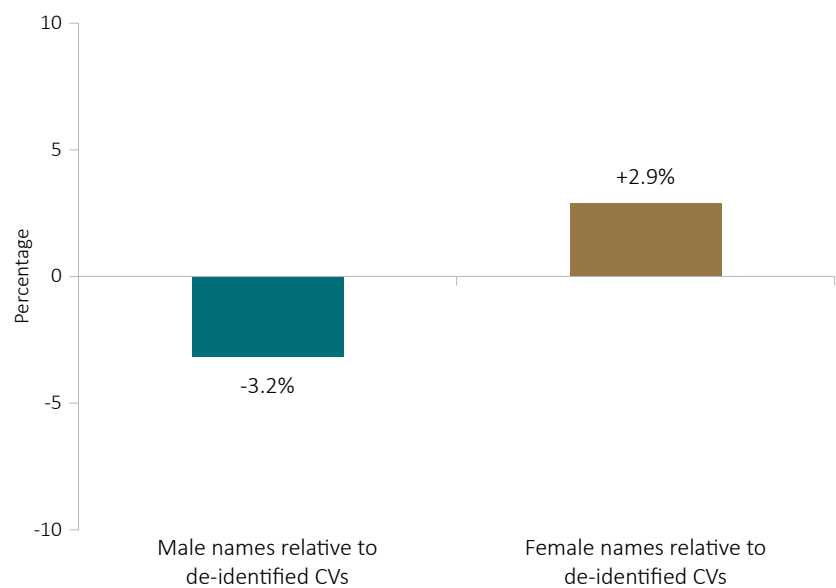
Note: Participants were randomised into three groups and asked to review CVs and shortlist the best five candidates. Participants in Control Group 1 saw the CVs including the names of each candidate. Participants in Control Group 2 saw the exact same CVs - but with one important difference: each name was substituted with a name of the opposite gender. Participants in the Treatment Group saw the exact same CVs but instead of names they were simply labelled CV1, CV2 and so on.

RESULTS OF THE TRIAL

Although the effect of de-identification is modest, it points to the existence of a form of subtle affirmative action taking place among reviewers. The public servants reviewing the job applicants engaged in discrimination that favoured female applicants and disadvantaged male candidates.

Figure 1 reports the key results on gender bias as a probability of being shortlisted. For a given set of CVs, assigning female identities increases the probability of the CV being shortlisted by 2.9% on average relative to the de-identified version. For the same set of CVs, assigning a male identity decreases the probability the CV is shortlisted by 3.2% on average. Both of these differences are small but are statistically significant at the 99% confidence level (meaning that the same result would be found 99% of the time if this exercise were repeated over and over again in the actual population). What this means is that, in practice, if applications were de-identified, we could expect that the likelihood of any female candidate being shortlisted would fall by 2.9%, on average, while likelihood of any male candidate being shortlisted would go up by 3.2%. Note that all results are reporting the probability of being shortlisted, not how the proportion of males and females in the shortlist composition changes with de-identification.


FIGURE 2: GENDER BIAS - WHAT IS THE EFFECT OF IDENTIFICATION ON THE SHORTLIST



Note: Candidates were shortlisted *more* when their names indicated they were female.
Male candidates were *less likely* to be shortlisted when their names were identifiable.

Figure 2 reports the results of the analysis of minority bias. Affirmative action towards the Indigenous female candidate is the largest, being 22.2% more likely to be shortlisted on average when identified compared to the de-identified condition. On the other hand, the identified Indigenous male CV is 9.4% more likely to be shortlisted on average compared to when it is de-identified. In absolute terms most minority candidates are on average more likely to be shortlisted when named compared to the de-identified condition, but the difference for the Indigenous female candidate is the only one that is statistically significant at the 95% confidence level.

FIGURE 3: MINORITY BIAS - WHAT IS THE EFFECT OF IDENTIFICATION ON THE SHORTLIST



When we examined specific sub-groups of APS staff we also discovered some interesting differences in behaviours. In particular:

- Overall, male reviewers displayed markedly more discrimination in favour of minority candidates than did female reviewers. Male reviewers were 11.6% more likely to shortlist minority men and 13.6% more likely to shortlist minority females, while female reviewers were only 1.84% more likely to shortlist minority men and 5.5% more likely to shortlist minority females, compared to the de-identified condition.
- APS staff aged 40+ displayed much stronger affirmative action in favour of female minorities than did staff under the age of 40. These reviewers were 10.0% more likely to shortlist minority females, while younger reviewers were only 5.8% more likely to shortlist female minorities, compared to the de-identified condition.
- APS staff working in human resources roles applied strong affirmative action in favour of both females and minorities: they were 9.0% more likely to shortlist females and 41.4% more likely to shortlist female minorities, compared to the de-identified condition.

- There was considerable variation in behaviour across agencies. For example, reviewers in some agencies appeared not to favour female or minority candidates to any significant extent. The agency displaying the strongest affirmative action for minority men was 55.4% more likely to shortlist minority men on average, when they could be identified, compared with when the candidates were de-identified.

The results from this trial demonstrate that, on the whole, public servants engage in positive discrimination towards female and minority candidates. De-identification of CVs in such a context has the effect of decreasing the number of female and minority candidates shortlisted for executive level APS positions.


To our knowledge, this study is the first of its kind to implement a rigorous trial to assess bias and measure the effect of de-identification in recruitment for a senior management position, certainly in the Australian (public or private sector) context. One previous Australian study pointed to discrimination disadvantaging ethnic minority applicants seeking entry-level employment (this involved sending fake CVs to job advertisements that did not require post-school qualifications, and subsequently measuring how call back rates varied across different ethnic groups)⁷. We find very different results when focussing on recruitment into executive level positions in the public service, where recruiters appear to be positively discriminating in favour of female and minority candidates.



LIMITATIONS OF THE TRIAL

There were some potential limitations to the study. In particular, as participation in this study was voluntary, it is possible that it attracted participants who are more likely to support diversity and gender equality. To address this issue, we gathered separate evidence via a survey conducted several weeks after the trial that asked a broad set of APS officers at the same levels their views on a range of issues. The survey evidence showed that attitudes towards diversity and gender equality among participants in our trial were representative of the broader population of employees at the same levels. The study sample was not distinctive or unusual in terms of their views on these issues.

Another important point to note is that, as this was a framed field experiment in which individuals knew they were part of a study, there is potential for subject reactivity or scrutiny bias. Even though this was a familiar task for participants, it is possible that they behaved differently than they would in a real recruitment situation. While we cannot control for the exercise being hypothetical, we were able to include design features that allowed us to mitigate reactivity and actually identify participants who might have guessed that the study was examining bias via their responses to some additional survey questions at the end of the exercise. When these few participants are excluded from the analysis, the results are unchanged. Our results would be further validated by a field trial using a real recruitment process to test the impact of de-identification of CVs on shortlisting.



POLICY LESSONS

Many organisations, including APS and other public sector agencies in Australia, are trialling the de-identification of job applications as a way of attempting to mitigate bias at the early stages of the recruitment process and promote diversity. Yet, the costs of de-identification can be high and the effects of de-identification are largely unknown. The existing evidence on the effects of de-identification is limited and mixed. Most importantly, the impact of de-identification can be expected to hinge critically upon the amount and direction of bias present in each agency. Given there were variations in the degree of affirmative action applied across sub groups and agencies in specific circumstances it may be appropriate to de-identify a selection process.

The overall implications of our study are that on average, across a broad range of APS agencies, introducing de-identification would have the unintended consequence of setting back efforts to promote more diversity at the senior management level in the public service. As things stand, senior public servants appear to be promoting diversity in the way they make decisions when selecting job candidates for shortlists during the initial stage of the recruitment process. This is not possible if applications are de-identified.

However, it remains clear that more work needs to be done to address the problem of gender inequality. The shortlisting stage of recruitment is one small piece of the diversity puzzle. Other stages of recruitment could be tested, including how positions are advertised, what information is requested from applicants, how interviews are conducted, and how hiring panels are selected and run. We also need to study factors that influence career trajectories once candidates are hired including opportunities for training and professional development, talent management, performance reviews, evaluations for promotions and flexible working arrangements. It is encouraging to see increased focus and activity in this space with many government and private sector organisations experimenting with different approaches and implementing new policies to attempt to improve diversity. Significant biases may be at work in each of these areas and eliminating or mitigating them will require innovative solutions and rigorous testing to discover what works.

Our results help to demonstrate the importance of testing interventions to address diversity before introducing them at full scale. An intervention that was thought to enhance the chances of individuals from traditionally disadvantaged groups being shortlisted for a senior role in the APS by removing bias, would have, in all likelihood, lessened their chances. The findings provide impetus for conducting more rigorous evaluations of new (and existing) initiatives aimed at countering explicit and implicit forms of discrimination and increasing gender, racial and ethnic diversity at

1. Department of Agriculture & Water Resources, Australian Taxation Office, Attorney General's Department, Department of Defence, Department of Foreign Affairs & Trade, Department of Employment, Department of Environment and Energy, Fair Work Ombudsman, Department of Health, Department of Industry, Innovation & Science, Offices of National Assessments, Department of the Prime Minister & Cabinet, Department of Social Services, Department of the Treasury
2. For a full report on the study, please see: Michael J. Hiscox and Lilia Arcos-Holzinger. 2017. "Going Blind to See More Clearly: The Effects of De-Identifying Job Applications in the Australian Public Service. BETA Working Paper 2017-1. Canberra.
3. Bertrand, M., Chugh, D., & Mullainathan, S. (2005). Implicit discrimination. *American Economic Review*, 94-98.
4. Goldin, C., & Rouse, C. (2000). Orchestrating impartiality: The impact of "Blind" Auditions on Female Musicians. *The American Economic Review*, 90(4), 715-741.
5. Norton, M. I., Vandello, J. A., & Darley, J. M. (2004). Casuistry and social category bias. *Journal of personality and social psychology*, 87(6), 817.
6. We acknowledge that Indigenous Australians are the traditional custodians and ancestors of Australia. As Indigenous Australians represent a small percentage of the overall population of Australia they can experience similar disadvantages to ethnic minorities. Hence, throughout the remainder of this report Indigenous Australians will be included in the analysis reported under 'minorities' whilst acknowledging that Indigenous cultures are complex and diverse.
7. Booth, A. L., Leigh, A., & Varganova, E. (2012). Does ethnic discrimination vary across minority groups? Evidence from a field experiment. *Oxford Bulletin of Economics and Statistics*, 74(4), 547-573.



