



Encouraging self-blinding in hiring

Sean Fath, Richard P. Larrick, & Jack B. Soll

abstract

One strategy for minimizing bias in hiring is *blinding*—purposefully limiting the information used when screening applicants to that which is directly relevant to the job and does not elicit bias based on race, gender, age, or other irrelevant characteristics. Blinding policies remain rare, however. An alternative to blinding policies is *self-blinding*, in which people performing hiring-related evaluations blind themselves to biasing information about applicants. Using a mock-hiring task, we tested ways to encourage self-blinding that take into consideration three variables likely to affect whether people self-blind: default effects on choices, people’s inability to assess their susceptibility to bias, and people’s tendency not to recognize the full range of information that can elicit that bias. Participants with hiring experience chose to receive or be blind to various pieces of information about applicants, some of which were potentially biasing. They selected potentially biasing information less often when asked to specify the applicant information they wanted to receive than when asked to specify the information they did not want to receive, when prescribing selections for other people than when making the selections for themselves, and when the information was obviously biasing than when it was less obviously so. On the basis of these findings, we propose a multipronged strategy that human resources leaders could use to enable and encourage hiring managers to self-blind when screening job applicants.

Fath, S., Larrick, R. P., & Soll, J. B. (2023). Encouraging self-blinding in hiring. *Behavioral Science & Policy*, 9(1), XX–XX. <https://doi.org/10.1353/XXXXX>

In a study published in 2017, researchers analyzed roughly five years of *pull requests*—that is, proposed changes to software projects—on the software development site GitHub to test whether requests made by women were evaluated differently than those made by men. The results were striking. When proposed changes came from software developers who were outsiders to a project (as opposed to project owners or known collaborators), project leaders were more likely to accept changes proposed by men than those proposed by women. However, this trend held only when the gender of the developer proposing the changes was identifiable. When project leaders were unable to discern the gender of the developer proposing the changes, they became more likely to accept proposals from women than from men.¹

This example demonstrates the value of a policy of *blinding*, or purposefully limiting the availability of irrelevant information that could potentially bias an evaluation of a person's ideas, qualifications, or performance. Blinding policies increase objectivity in evaluations by preventing evaluators from receiving information that might bias their assessments. In the domain of work, a hiring manager who does not know the name of a job applicant—say, because the name has been stripped from the applicant's resume—cannot possibly use that name to make assumptions about the applicant's race, gender, or other attributes peripheral to job performance. Stereotypes about race and gender cannot then leak into assessments of other information, such as job credentials.²

Yet when it comes to making hiring decisions, blinding policies remain relatively rare. Although a handful of boutique firms, such as GapJumpers and Applied, have emerged to help companies perform blind initial screens of job applicants, we have found that few institutions choose to use such services or establish internal blinding policies for the hiring process. In a survey we reported on in 2021, we asked more than 800 human resources (HR) professionals—who averaged 14 years of experience in the field—about whether they had experience with blinding policies in the hiring process.³ We

found that 81% of them had never worked at an organization that used blinding policies at any point during hiring. Moreover, 80% indicated they had never received training about blinding as a possible bias-reduction strategy.

Although these data are not representative of all U.S. organizations, they suggest that blinding policies and services are not commonly used in hiring. Some other alternative hiring practices, such as artificial intelligence–based screens of applicants, may be considered blind to the extent they are machine-driven, but these practices can still result in biased evaluations. For instance, automatically screening out applicants who have gaps in employment can affect women disproportionately, because women have employment gaps more often than men do.^{4,5} Further, machine-driven practices are often used in conjunction with nonblind human evaluations.^{6–8}

Blinding may be uncommon in institutional hiring in part because the jobs of most organizations vary widely in qualifications and duties. As a result, HR professionals may be concerned that uniform rules may not be appropriate in all cases. In addition, they may want to avoid limiting the autonomy of managers in making hiring decisions, as hiring managers tend to balk at initiatives such as diversity-fostering hiring policies that limit their latitude in decision-making.⁹ Yet without blinding, hiring decisions may be compromised by bias from information that is not directly related to job qualifications. Biasing information—such as a person's name, age, or appearance—is often either included in applicant materials¹⁰ or easily gathered from the internet.^{11,12}

In this article, we ask, is it possible to encourage those making hiring-related decisions to self-blind—to choose on their own not to receive biasing information about applicants? Encouraging self-blinding during the initial screening of applicants would preserve hiring managers' autonomy as well as the flexibility needed to adapt the hiring process to particular jobs. For instance, an organization could introduce a checklist-based system by which a manager could pick which information to see or not to see

when evaluating job candidates. Such a system represents a behavioral *nudge*—a gentle push to do something that does not limit autonomy or choices. It could prompt managers to avoid seeing biasing information without limiting their freedom, thereby reducing employment discrimination based on race, age, gender, or any other job-irrelevant characteristic.

To address our question, we explored the influence of three key factors on whether hiring managers performing an initial screen of applicants would blind themselves to biasing information about those applicants: the psychological pull of defaults, people's sense of their susceptibility to bias, and people's understanding of what information can lead to bias. We examined the effects of each of these factors in a mock hiring task, with the overall aim of determining the most effective design for a self-blinding process in organizations. Next, we discuss the science undergirding our exploration of these factors and our predictions about their effects.

Factors Influencing the Likelihood of Self-Blinding

Default Effects

Hiring decisions are typically structured such that hiring managers receive biasing information about job applicants by default. For instance, hiring managers often learn applicants' names at the beginning of the hiring process, and a name may provide information such as a person's race, gender, and social class. Biases about these social categories can then distort the way the manager evaluates the person's suitability for the job. To avoid this distortion, a hiring manager could choose to remain unaware of applicants' names, but that scenario is unlikely if managers get this information by default. The literature on default effects shows that decision-makers in many domains tend to accept defaults.

For instance, employees are more likely to participate in a retirement savings plan when their employer enrolls them by default, relative to when the default state is nonenrollment and participation requires employees to make an

effort to sign up, or opt in, to the plan.¹³ People are more likely to be organ donors,¹⁴ undergo HIV screening,¹⁵ and get the flu vaccine¹⁶ when arrangements are made for them (forcing them to opt out to avoid participation) than when they must opt in to participate. (See the Supplemental Material for more information about default effects.)

Similarly, if hiring managers receive all information—including biasing information—about an applicant by default, they may be disinclined to depart from that default state to avoid receiving the biasing information. We therefore predicted that providing no information unless items were specifically requested (that is, unless managers opted in to receiving particular items) would result in managers being more likely to blind themselves to biasing information than would providing all information and requiring managers to opt out of seeing particular items. Put another way, we hypothesized that using an opt-in framework would be the optimal strategy for nudging hiring managers to blind themselves to biasing information.¹⁷

Perceived Susceptibility to Bias

Hiring managers' inclination to self-blind to biasing information about job applicants may also be shaped by their personal sense of susceptibility to bias in hiring decisions. Unfortunately, people are poor judges of their own propensity for bias. In social psychological research on self-perceived bias susceptibility, participants often judge themselves to be objective in their specific evaluations¹⁸ and general perceptions of the world¹⁹ and believe they are less susceptible to bias than others are.²⁰ This misperception may make hiring managers more likely to elect to see biasing information for themselves than they would be if they were making the choice for someone else. To test this proposition, we asked some of our participants to consider what choice they would make for others regarding whether they should see biasing information. We predicted that these participants would be more likely to avoid providing biasing information to others doing the screening than would participants instructed to make that choice for themselves. If that prediction proved correct, the finding

would indicate that people’s misperception of their own susceptibility to bias at least partly affects whether they choose to look at biasing information. This misperception might be counteracted by asking managers to make a choice for someone else before choosing options for themselves.

Bias Transparency

The third factor that can influence whether hiring managers blind themselves to biasing information involves the nature of that information and whether managers recognize its potential to bias decisions. Some of the information that can bias decisions about job applicants may not be obviously biasing. For instance, an applicant’s name may appear to be innocent background information even though it may indicate a person’s gender and race, among other attributes. By contrast, explicit mention of a person’s gender or race is transparently biasing. We tested how often participants chose to see transparently biasing versus nontransparently biasing information. We anticipated that participants would consider nontransparently biasing information to be less biasing than more overtly biasing information and thus would elect to see nontransparently biasing information more often than they would choose to see overtly biasing information. If so, this propensity would need to be considered when strategies nudging self-blinding are designed.

The Predictions in Brief

In a nutshell, we predicted that participants would be more likely to blind themselves to potentially biasing information (whether transparently or nontransparently biasing) when they had to opt in (specifically choosing what information to see) than when they had to opt out of receiving the information. We also predicted that participants would choose potentially biasing information for review less often when making the choice for others than when making it for themselves. Finally, we predicted that in any of those conditions, participants would elect to see transparently biasing information less often than they would elect to see nontransparently biasing information, even though both types could, in fact, bias their decisions.

We also tested whether self-blinding nudges might affect participants’ interest in seeing information that is important for making a good decision. Strategies to encourage bias reduction by self-blinding should be adopted only if they do not markedly suppress hiring managers’ inclination to receive useful information about job applicants—that is, information relevant to applicants’ job qualifications. We did not expect self-blinding nudges to inhibit participants from electing to see information that is widely accepted to be diagnostic of job performance, because this information is not likely to be viewed as a source of bias. That is, we expected that participants would be just as likely to ask to see useful information regardless of the decision-making frame (opt in or opt out) or whether they were making the decision for themselves or for others.

Method

We recruited 800 participants with hiring experience to take part in our experiment, targeting about 100 participants for each of the eight study conditions we planned; we received 798 complete responses.²¹ The mean age of the participants was 39.82 years ($SD = 12.03$); 47.4% were women. Participants had an average of 19.16 years of work experience and estimated that they had made an average of 36.67 hiring decisions in their careers. They were all U.S. citizens and were recruited through an online platform (<https://www.prolific.co/>) that supplies research participants.

Participants completed a mock hiring task in which they screened applicants for a hypothetical position at their place of work to determine whom to advance to the interview stage. All participants received a checklist from which they could choose to see any of seven types of information available about applicants. Five of the seven items on the checklist represented useful information, which we define as information that is commonly accepted to be relevant for hiring decisions. We selected these items—the job applicant’s college, major, previous work experience, job-related skills, and references—using a pool of sample applications for U.S. jobs posted online as a guide.

The remaining two items on our checklist were those that we prejudged to be irrelevant to job performance and potentially biasing. Participants saw one of two sets of items, depending on their study condition. The first set of two items consisted of a job applicant's race and gender, which we deemed to be transparently biasing. The second set of two items consisted of a job applicant's picture and name, which we judged to be nontransparently biasing. All items were presented to participants in a randomized order. (All materials and data for our study are archived online at <https://osf.io/2vthn/>.)

To assess the effects of an opt-out or opt-in framework on self-blinding preferences, we randomly assigned participants to one of two sets of instructions: One told participants to tick the boxes next to the items they did not wish to receive (that is, to opt out of the default of receiving all the information), and the other told them to tick the boxes next to the items they wanted to receive (that is, to opt in to receiving specific information). To assess self-perceived susceptibility to bias, we randomly assigned the participants in the opt-out and opt-in conditions to either choose the information they wanted to receive if they were making the screening decision themselves or decide what information to provide to someone else doing the screening. Finally, we further divided those four groups, randomly assigning participants to use a checklist that included either the two transparently biasing items or the two nontransparently biasing items. For each of the resulting eight

conditions (see Table 1), we tabulated the items participants chose to see.

To confirm that participants agreed with us on which items were useful versus biasing in relation to a hiring decision, we conducted a posttest using a separate group of 104 participants with hiring experience. The results generally supported our classifications of these items as useful (five items) or potentially biasing (two items). One exception was the name of the job applicant's college, which we had prejudged to be useful but posttest participants rated as slightly more biasing than useful. (See the Supplemental Material for details.) As a result, we did not use the name of the job applicant's college item in the analyses that follow.

Results

Overall, our hypotheses were supported. Participants were less likely to choose to see the biasing information (name, picture, gender, and race) when they were instructed to opt in to information they wanted to see ($M = 22.3\%$) than when they had to opt out to exclude information they did not want to see ($M = 32.1\%$, $p < .001$). Participants were also less likely to choose information that was potentially biasing when making a choice for others ($M = 21.4\%$) than for themselves ($M = 32.9\%$, $p < .001$). Finally, participants were less likely to elect to see biasing information when the possibility of bias was relatively transparent, as in a person's race or gender ($M = 17\%$) than when it was

Table 1. The eight conditions studied, by mix of frame, person receiving the information, & type of potentially biasing information available

Frame	Person receiving the information	Type of potentially biasing information
Opt out	Self	Transparently biasing
Opt out	Self	Nontransparently biasing
Opt out	Others	Transparently biasing
Opt out	Others	Nontransparently biasing
Opt in	Self	Transparently biasing
Opt in	Self	Nontransparently biasing
Opt in	Others	Transparently biasing
Opt in	Others	Nontransparently biasing

Note. The opt-out frame involves choosing what to exclude. The opt-in frame involves choosing what to include.

nontransparent, as in a person's picture or name ($M = 37.3\%$, $p < .001$). These main effects are averages across all of the conditions. (See the Supplemental Material for more details of our analyses and results.)

Next, we compared the effects of the opt-in versus opt-out frames and self versus other decisions on choices to receive biasing versus useful applicant information. Across conditions, the vast majority of participants asked to see the useful information ($M = 90.9\%$), whereas only about a quarter of the participants asked to see the biasing information ($M = 27.2\%$, $ps < .001$).

When we looked at the effects of opting out or opting in specifically on the selection of transparently or nontransparently biasing information, we found that participants were less likely to choose to see the nontransparently biasing items (a picture or name) when they were in an opt-in condition in which they actively chose to see items ($M = 31.3\%$) than when they were in an opt-out condition in which they excluded items from a list ($M = 43.4\%$, $p = .002$). The same pattern held for the transparently biasing items (race and gender): A mean of 21.1% of the participants in an opt-out condition but a mean of only 12.8% of the participants in an opt-in condition selected the transparently biasing items ($p = .015$). However, for the useful items, the opt-out versus opt-in distinction had a much smaller effect, such that roughly 9 out of 10 participants chose to see the useful information regardless of the default frame. (The mean percentage of the opt-out conditions was 92.5%; that of the opt-in conditions was 89.4%, $p = .026$.)

Similarly, participants were less likely to choose the nontransparently biasing information for others ($M = 30.2\%$) than for themselves ($M = 44.1\%$, $p < .001$), and the same was true for the transparently biasing items: A mean of 12.9% of participants in the other conditions chose these items compared with a mean of 21.3% in the self conditions ($p = .015$). In comparison, participants selected the useful items for themselves at roughly the same rate as they did for others, with about 9 out of 10 choosing the

information in either case: A mean of 90.6% in the self conditions and a mean of 91.2% in the other conditions ($p = .655$).

The panels of Figure 1 break down the data further, showing the percentage of participants who chose for themselves (Panel A) or for others (Panel B) each of the useful and biasing items that were available, broken down by opt-out and opt-in conditions. Figure 2 shows the percentage of participants in each condition who chose useful information in aggregate (job skills, work experience, references, and college major) and biasing information in aggregate (name, picture, gender, and race).

Discussion


In our study, participants with hiring experience blinded themselves to information that was potentially biasing about mock job applicants more often when (a) they needed to opt in to see information about the applicants than when they had to opt out, (b) they were making a decision for someone else rather than for themselves, and (c) the biasing items were transparently biasing (such as the item identifying race) rather than more subtly biasing (such as the item providing a name). Next, we discuss ways that companies and other institutions may leverage these findings to encourage hiring managers and others making hiring-related decisions to blind themselves to potentially biasing information about job applicants.

Solutions

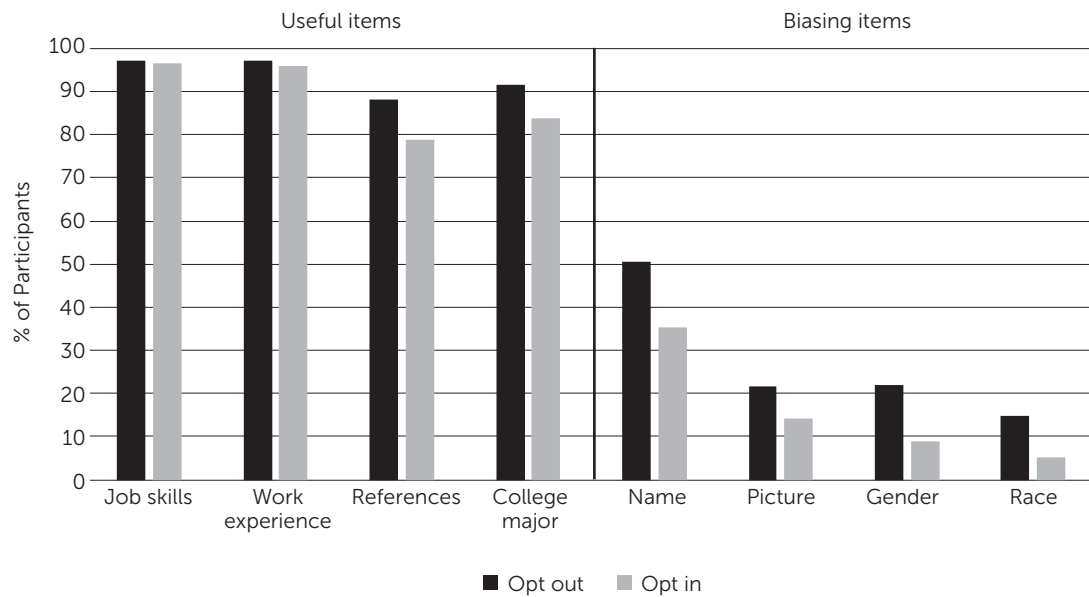
Leverage Default Effects. In our study, the opt-out scenario created a default in which participants would receive all the information on a checklist unless they opted out of some of it. In the opt-in scenario, the default was receiving no information. Research on default effects has predominantly demonstrated that people are more likely to adopt beneficial policies or behaviors under opt-out conditions than opt-in ones, such as when a person who does not want to be an organ donor has to opt out when obtaining a driver's license. In this study, however, the opt-in condition was more effective at minimizing the selection of biasing

Figure 1. Percentages of participants choosing for themselves & for others whether to see applicant information

A. Participants choosing information for themselves




B. Participants choosing information for others



Note. When making a choice for themselves (A), participants generally asked to see biasing applicant information (four items at right) less often when they were given a list of possible items and asked to specify the ones they wanted to see (that is, when they were in an opt-in condition) than when they were instructed to specify which ones they did not want to see (that is, when they were in an opt-out condition). The same was true for participants who made choices for others (B). The same was not true for useful—that is, clearly job-relevant—information. In addition, whether choosing for themselves or for others, significantly fewer participants selected transparently biasing information (gender and race) than nontransparently biasing information (name and picture). Participants who chose information for others were less likely to seek biasing information than were those who chose information for themselves.

Figure 2. Percentages of participants choosing useful & biasing information, by condition



Note. Participants chose biasing information (name, picture, gender, and/or race) least often when they were in an opt-in condition and were choosing for others and most often when they were in an opt-out condition and choosing for themselves. The differences between conditions were attenuated for useful information.

information and did so without markedly diminishing interest in useful information relevant to job performance. In certain domains, such as hiring, some options (such as seeing job-related skills) are likely to be favored regardless of whether the option is provided by default.

So although most research on default effects underscores the effectiveness of interventions that allow people to make passive decisions,^{13,17} such as sticking with a desirable default, our work suggests that requiring active decision-making is best for nudging managers to self-blind to biasing information. In our study, participants who had to opt in to see information (that is, who made an active decision to look at each item) seemingly became more attentive to which items might bias their decisions and consequentially became less likely to select items providing biasing information. This take-away is consistent with research demonstrating that inclusion frames (which require people to choose the best items from a broader list) foster more deliberative thinking than do exclusion frames (which require people to reject the worst items from a broader list).¹⁷ The findings are also consistent with research showing that when choices to receive biasing information are

driven by curiosity, curiosity-driven impulses can be reduced by using decision frames that cue deliberative reasoning.²²

Our results suggest that organizations could nudge hiring managers to selectively self-blind to biasing information by instituting a checklist system in which managers must pick the information they wish to see about applicants. An organization could have a dedicated employee create the checklist by itemizing the information available about applicants and then give that list to the hiring managers. Such a process may also be appealing to decision-makers, who tend to prefer opt-in to opt-out frames when making choices.²³

Circumvent the Bias Blind Spot. Our finding that participants selected biasing information for themselves more often than they did for others is consistent with other research showing that people perceive others to be more susceptible to bias than they themselves are.²⁰ In our study, this difference was attenuated when the information was patently useful and relevant. These results suggest that to encourage hiring managers to self-blind, organizations could train hiring managers to consider what information

they would give to someone else making a hiring decision before making that choice for themselves. Because hiring managers are likely to want their decisions to be consistent, they are likely to make the same choices for themselves as they did for others.²⁴ A training module encouraging hiring managers to “consider what information you would want someone else making this decision to have” could be included in organizations’ antibias training and would likely offer benefits beyond encouraging self-blinding in hiring.

Boost Awareness of Hidden Bias. Our finding that participants selected nontransparently biasing information at a higher rate than they selected transparently biasing information is noteworthy because much of the biasing information available to hiring managers may be nontransparently biasing. For instance, applicants’ names are commonly provided on applications, and photos are typically available on applicants’ social media pages or personal websites, which are information sources managers often use.^{11,12,25} Although a name or a picture may be less obviously biasing than a person’s noted gender or race, they are often just as biasing or more so. Both often convey race and gender, and a photo is likely to communicate additional biasing information such as attractiveness, age, and physical fitness. Many studies have documented hiring bias related to applicants’ names.¹⁰

Yet our results show that many people are not aware of how biasing names and photos can be. Similarly, people might not realize that being aware of applicants’ college graduation years may trigger biases related to age, that knowing applicants’ hobbies could lead to biases related to social class or disability status,²⁶ and so on. Even information that is objectively nonbiasing and merely irrelevant to a hiring decision is a good candidate for blinding, because, at best, the inclusion of such information adds noise to evaluations.²⁷ Our results suggest that hiring managers would be more likely to self-blind to biasing information the more they are made aware—perhaps through continuing education—of the potentially biasing or at least noise-inducing content lurking within

seemingly innocent information. To ensure that hiring managers get this education, organizations could require them to complete a training module on hidden sources of bias in hiring-related information.

Combat Belief in the Usefulness of Biasing Information. We understand that which applicant information is biasing versus relevant is likely to vary across industries and that even the information deemed biasing in our experiments may be directly relevant in some contexts. For instance, a photo is likely to be useful and relevant for a modeling job.

But hiring managers often believe that potentially biasing information is useful when it is not. In our study, for instance, participants may have sought to derive information about cultural fit from applicants’ photos, as we have observed in other studies.²² It is also possible that some participants chose to view potentially biasing information because they wanted to favor applicants from marginalized groups. Although well-intentioned, this type of choice carries dangers, because removing unwanted bias from one’s reasoning is difficult.² As an example, a well-intentioned hiring manager might seek out an applicant’s photo to clarify their race or gender, with the goal of favoring applicants from marginalized groups. But in the process, that manager opens the door to biases related to age and attractiveness. Corporate training sessions on bias can help hiring managers understand that biasing information is typically more harmful than helpful, which, in turn, should increase their preference for self-blinding.

Other Considerations Related to Self-Blinding

Could self-blinding ever be counterproductive to combating discrimination? As we have noted, hiring managers may want to see information that will help them favor applicants from disadvantaged social groups. Another argument against blinding is that knowing applicants’ social group status might provide important context for assessing their credentials. If, for example, a person had to overcome a lifetime of disadvantage or discrimination to gain those

credentials, would it not be helpful to take that past into account? Doing so might help minimize the advantages of members of dominant groups whose qualifications might derive in part from privilege. It is generally for these reasons that the merits of “colorblind” policies are questioned.^{28,29}

However, the bulk of the field studies assessing the effects of using blind initial screens suggest that members of marginalized groups, such as women and ethnic minorities, are often more likely to reach the interview stage when initial screens are anonymized.^{30–32} Moreover, multiple recent reviews and meta-analyses show that unblinded initial screens of applicants decrease the likelihood of members of marginalized groups receiving callbacks for interviews.^{10,33–36} These findings strongly indicate that a blinding process during applicant screening can, via a reduction in discrimination, help achieve an institution’s goal of diversity in hiring.

Moreover, self-blinding in applicant screening may have carryover benefits in interviews. If the

same evaluator who performed an initial blind screen also interviews the selected job candidates, that evaluator is likely to continue to try to discount biasing demographic information to maintain a consistent strategy throughout the process.^{24,37}

Still, self-blinding is just one tool among many that may be used to achieve diversity goals in hiring. Self-blinding nudges should be used in tandem with other strategies, such as unblind targeted recruiting to increase the proportion of people from marginalized groups who apply for a job in the first place—for instance, by establishing talent development or pipeline programs at historically Black colleges and universities—and structured interviewing procedures that decrease the likelihood of bias against members of marginalized groups in face-to-face interviews. A multifaceted strategy is necessary to address bias in hiring decisions, and self-blinding is one important component of that strategy. (See the Supplemental Material for a fuller discussion of hiring procedures that promote diversity in hiring.)

In Brief: How Organizations Can Encourage Hiring Managers to Self-Blind

Human resources professionals can decrease the chances of bias in hiring practices in their organizations by enabling and encouraging hiring managers to blind themselves to biasing information about applicants during the initial screening process. Here are the actions we recommend, based on our research:

Appoint an intermediary to create a checklist. Assign a dedicated employee to read applications, categorize the information in them, and make a checklist for the hiring manager listing all the available types of information about applicants—for example, work experience, college major, references, name, and photograph.

Draft instructions for checklist use that encourage self-blinding.

- Ask hiring managers to first consider what information they would want another hiring manager to see. For example, “Imagine a situation in which someone else is tasked with hiring someone for an open position at your organization, and it is up to you to decide what information they incorporate into their decision.” Then ask the manager to pick that information for themselves.
- Ask managers to pick the information they want to see rather than asking them to pick the information they do not want to see (that is, use opt-in rather than opt-out framing). For example, “Here are the pieces of information about applicants that are available. Please tick the box(es) next to the information you want to see.”

Offer or require training for managers about hidden sources of bias and the importance of blinding themselves to biasing information. This training should help managers spot the potential for bias in information that does not necessarily seem like it would trigger biases, such as a person’s name or photograph. Training could also help combat beliefs in the potential usefulness of biasing information.

Conclusion

Until blinding policies become commonplace in hiring, seemingly innocuous information about job applicants, such as their name, hobbies, or college graduation year, will continue to enable discrimination. In this article, we have discussed self-blinding, in which hiring managers choose on their own to avoid information about applicants that could bias or distort their evaluations, and we have identified three factors that could influence whether managers self-blind.

Our research suggests steps that organizational leaders could take to encourage hiring managers to self-blind when screening job applicants in the early stages of the hiring process (see the sidebar *In Brief: How Organizations Can Encourage Hiring Managers to Self-Blind*). We propose that organizations appoint a dedicated employee to itemize the available information about applicants. Evaluators could then be instructed to think about which types of information they would provide to a peer (because people tend to be stricter when choosing for others than for themselves) before opting in to receiving the information for themselves. Organizations should also provide training concerning the types of information that carry bias, myths about the usefulness of such information, and how to combat misperceptions of one's own susceptibility to bias. We believe that nudging self-blinding using the principles and guidelines outlined here will result in fairer and more accurate decisions in hiring.

author affiliation

Fath: Cornell University. Larrick and Soll: Duke University. Corresponding author's e-mail: sean.fath@cornell.edu.

supplemental material

- <http://behavioralpolicy.org/journal>
- Method & Analysis

references

1. Terrell, J., Kofink, A., Middleton, J., Rainear, C., Murphy-Hill, E., Parnin, C., & Stallings, J. (2017). Gender differences and bias in open source: Pull request acceptance of women versus men. *PeerJ Computer Science*, 3, Article e111. <https://doi.org/10.7717/peerj-cs.111>
2. Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin*, 116(1), 117–142. <https://doi.org/10.1037/0033-2909.116.1.117>
3. Fath, S., & Zhu, S. (2021). *Preferences for, and familiarity with, blinding among HR practitioners*. SSRN. <https://doi.org/10.2139/ssrn.3768039>
4. Ajunwa, I. (2019). The paradox of automation as anti-bias intervention. *Cardozo Law Review*, 41(5), 1671–1742. <http://cardozolawreview.com/the-paradox-of-automation-as-anti-bias-intervention/>
5. Ajunwa, I., Friedler, S., Scheidegger, C. E., & Venkatasubramanian, S. (2016). *Hiring by algorithm: Predicting and preventing disparate impact*. SSRN. <http://sorelle.friedler.net/papers/SSRN-id2746078.pdf>
6. Ajunwa, I. (2019, October 8). Beware of automated hiring. *The New York Times*. <https://www.nytimes.com/2019/10/08/opinion/ai-hiring-discrimination.html>
7. Teichman, D., Talley, E., Egidy, S., Engel, C., Gummadi, K., Hagel, K., Lewandowsky, S., MacCoun, R., Utz, S., & Zamir, E. (2021). Institutions promoting or countering deliberate ignorance. In R. Hertwig & C. Engel (Eds.), *Deliberate ignorance: Choosing not to know* (pp. 275–298). MIT Press. <https://doi.org/10.7551/mitpress/13757.003.0022>
8. Lewis, N. (2018). *Will AI remove hiring bias?* Society for Human Resource Management. https://www.shrm.org/LearningAndCareer/learning/educational-program-materials/Documents/SHRM_Will%20AI%20Remove%20Hiring%20Bias.pdf
9. Dobbin, F., Schrage, D., & Kalev, A. (2015). Rage against the iron cage: The varied effects of bureaucratic personnel reforms on diversity. *American Sociological Review*, 80(5), 1014–1044. <https://doi.org/10.1177/0003122415596416>
10. Bertrand, M., & Duflo, E. (2017). Field experiments on discrimination. In A. V. Banerjee & E. Duflo (Eds.), *Handbook of economic field experiments* (Vol. 1, pp. 309–393). Elsevier. <https://doi.org/10.1016/bs.hefe.2016.08.004>
11. Acquisti, A., & Fong, C. (2020). An experiment in hiring discrimination via online social networks. *Management Science*, 66(3), 1005–1024. <https://doi.org/10.1287/mnsc.2018.3269>
12. Bartoš, V., Bauer, M., Chytilová, J., & Matejka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, 106(6), 1437–1475. <https://doi.org/10.1257/aer.20140571>
13. Madrian, B. C., & Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics*, 116(4), 1149–1187. <https://doi.org/10.1162/003355301753265543>
14. Johnson, E. J., & Goldstein, D. (2003, November 21). Do defaults save lives? *Science*, 302(5649), 1338–1339. <https://doi.org/10.1126/science.1091721>
15. Montoy, J. C. C., Dow, W. H., & Kaplan, B. C. (2016). Patient choice in opt-in, active choice, and opt-out HIV screening: Randomized clinical trial. *The BMJ*, 352, Article h6895. <https://doi.org/10.1136/bmj.h6895>
16. Chapman, G. B., Li, M., Colby, H., & Yoon, H. (2010). Opting in vs opting out of influenza vaccination. *JAMA*, 304(1), 43–44. <https://doi.org/10.1001/jama.2010.892>
17. Levin, I. P., Huneke, M. E., & Jasper, J. D. (2000). Information processing at successive stages of decision making: Need for cognition and inclusion–exclusion effects. *Organizational Behavior and Human Decision Processes*, 82(2), 171–193. <https://doi.org/10.1006/obhd.2000.2881>
18. Uhlmann, E. L., & Cohen, G. L. (2007). “I think it, therefore it’s true”: Effects of self-perceived objectivity on hiring discrimination. *Organizational Behavior and Human Decision Processes*, 104(2), 207–223. <https://doi.org/10.1016/j.obhdp.2007.07.001>
19. Armor, D. A. (1999). *The illusion of objectivity: Bias in the belief in freedom from bias* [Unpublished doctoral dissertation]. University of California, Los Angeles.
20. Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the eye of the beholder: Divergent perceptions of bias in self versus others. *Psychological Review*, 111(3), 781–799. <https://doi.org/10.1037/0033-295X.111.3.781>
21. Gervais, W. M., Jewell, J. A., Najle, M. B., & Ng, B. K. (2015). A powerful nudge? Presenting calculable consequences of underpowered research shifts incentives toward adequately powered designs. *Social Psychological and Personality Science*, 6(7), 847–854. <https://doi.org/10.1177/1948550615584199>
22. Fath, S. B., Larrick, R. P., & Soll, J. (2022). Blinding curiosity: Exploring preferences for “blinding” one’s own judgment. *Organizational Behavior and Human Decision Processes*, 170, Article 104135. <https://doi.org/10.1016/j.obhdp.2022.104135>
23. Goodman, J. K., & Reczek, R. W. (2021). Choosing what to choose from: Preference for inclusion over exclusion when constructing consideration sets from large choice sets. *Journal of Behavioral Decision Making*, 34(1), 85–98. <https://doi.org/10.1002/bdm.2199>
24. Festinger, L. (1957). *A theory of cognitive dissonance* (Vol. 2). Stanford University Press.
25. Jobvite. (2014). *2014 Social Recruiting Survey* [Slides]. https://www.jobvite.com/wp-content/uploads/2014/10/Jobvite_SocialRecruiting_Survey2014.pdf
26. Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review*, 77(6), 999–1022. <https://doi.org/10.1177/0003122412463213>
27. Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, 11(3), 359–372. <https://doi.org/10.1177/1745691616635594>
28. Markus, H. R., Steele, C. M., & Steele, D. M. (2000). Colorblindness as a barrier to inclusion: Assimilation and nonimmigrant minorities. *Daedalus*, 129(4), 233–259.
29. Rattan, A., & Ambady, N. (2013). Diversity ideologies and intergroup relations: An examination of color-blindness and multiculturalism. *European Journal of Social Psychology*, 43(1), 12–21. <https://doi.org/10.1002/ejsp.1892>
30. Åslund, O., & Skans, O. N. (2012). Do anonymous job application procedures level the playing field? *ILR Review*, 65(1), 82–107. <https://doi.org/10.1177/001979391206500105>
31. Bøg, M., & Kranendonk, E. (2011). *Labor market discrimination of minorities? Yes, but not in job offers* (No. 33332). Munich Personal RePEc Archive. https://mpra.ub.uni-muenchen.de/33332/1/MPRA_paper_33332.pdf

32. Krause, A., Rinne, U., & Zimmermann, K. F. (2012). Anonymous job applications in Europe. *IZA Journal of European Labor Studies*, 1, Article 5. <https://doi.org/10.1186/2193-9012-1-5>
33. Quillian, L., Heath, A., Pager, D., Midtbøen, A. H., Fleischmann, F., & Hexel, O. (2019). Do some countries discriminate more than others? Evidence from 97 field experiments of racial discrimination in hiring. *Sociological Science*, 6, 467–496. <https://doi.org/10.15195/v6.a18>
34. Quillian, L., Pager, D., Hexel, O., & Midtbøen, A. H. (2017). Meta-analysis of field experiments shows no change in racial discrimination in hiring over time. *Proceedings of the National Academy of Sciences, USA*, 114(41), 10870–10875. <https://doi.org/10.1073/pnas.1706255114>
35. Zschirnt, E., & Ruedin, D. (2016). Ethnic discrimination in hiring decisions: A meta-analysis of correspondence tests 1990–2015. *Journal of Ethnic and Migration Studies*, 42(7), 1115–1134. <https://doi.org/10.1080/1369183X.2015.1133279>
36. Pager, D., & Shepherd, H. (2008). The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets. *Annual Review of Sociology*, 34, 181–209. <https://doi.org/10.1146/annurev.soc.33.040406.131740>
37. Stone, J., & Fernandez, N. C. (2008). To practice what we preach: The use of hypocrisy and cognitive dissonance to motivate behavior change. *Social and Personality Psychology Compass*, 2(2), 1024–1051. <https://doi.org/10.1111/j.1751-9004.2008.00088.x>