Journal of Experimental Social Psychology 114 (2024) 104628

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Beyond first impressions: Investigating the influence of visual attention and cue availability in discriminatory behavior[☆]

Eliane Roy [*], Y. Doug Dong, A. Ross Otto, Jordan Axt

McGill University, Canada

ARTICLE INFO

Keywords:
Discrimination
Social judgment
bias
Attractiveness
Eye-tracking

ABSTRACT

In many contexts, the magnitude of discrimination in social judgment is determined by the level of sensitivity and bias in evaluation. However, little is known about factors that shape these processes. Using a mock admissions task, we investigated how variation in the time spent processing non-diagnostic social information (e.g., a face communicating attractiveness) versus decision-relevant information (e.g., information about candidate's qualifications) differentially impacted sensitivity versus bias, using both correlational eye-tracking (Study 1) and experimental approaches (Studies 2–3). Higher sensitivity (i.e., less judgment errors) was consistently related to the amount of time participants viewed decision-relevant information. However, bias (i.e., selection leniency based on attractiveness) was unrelated to the amount of time participants chose to view or were allowed to view faces. Bias emerged when faces were shown for as little as 350 milliseconds. The ease with which social information is encoded suggests that merely limiting its' presentation is ineffective for reducing discrimination.

Decades of research have found that discrimination based on social information – like demographic information or physical characteristics – contributes to group disparities in outcomes like admissions or hiring (Milkman, Akinola, & Chugh, 2012; Rooth, 2009). Given these well-documented disparities, research in psychology has sought to better understand the processes that contribute to discrimination, as identifying inputs that do (or do not) contribute to socially biased judgment may facilitate the development of interventions that can reduce discrimination.

However, while prior work has made great strides in identifying the targets or contexts where discrimination is likely to emerge, the field has struggled to find particularly strong *psychological* predictors of discriminatory behavior. For example, recent meta-analyses found only small correlations between discriminatory behavior and implicit attitudes ($r < 0.20$; Forscher et al., 2019; Kurdi et al., 2019). Similar results have emerged when looking at self-reported desire to control prejudice (aggregate $r = 0.09$ in Axt & Lai, 2019), dehumanization (maximum $r = 0.23$ in Bruneau, Szekeres, Kteily, Tropp, & Kende, 2020), perceived intergroup threat (maximum $r = 0.27$ in Kauff & Wagner, 2012) or self-serving motives (meta-analytic $r = 0.09$ in Munder, Becker, & Christ, 2020). Such findings suggest that researchers may need to look to additional processes when seeking to identify factors that predict discrimination. Building on this earlier research, the present work investigated the potential role of a more proximal, yet comparatively underexplored predictor in discriminatory behavior: cue-specific attention.

Prior work in person perception has found that social information, such as race or gender, can be gleaned from target faces quickly and effortlessly (Ito & Urland, 2003; Zhou et al., 2020). Even more abstract traits like trustworthiness and competence can be predicted with some accuracy within 100 ms of exposure to individual targets (Todorov, Pakrashi, & Oosterhof, 2009; Willis & Todorov, 2006), and gleaned from groups within 500 ms (Chwe & Freeman, 2023). Together, this work has been crucial in illustrating just how automatically our perceptual system detects various forms of social information. However, one limitation of these studies is that participants were – either directly or indirectly – cued to using facial information to guide interpersonal judgment. That is, in these prior studies, participants were either 1) encouraged to use facial information in judgment (e.g., through instructions encouraging a reliance on first impressions; Cañadas, Rodríguez-Bailón, Milliken, & Lupiáñez, 2013), 2) told beforehand about the social dimension being investigated (e.g., monoracial or biracial status in Chen & Hamilton, 2012), 3) cued to the relevant social dimension when asked to report on it after the first face presentation (e.g., racial prototypicality in Dunham,

---

[☆] This paper has been recommended for acceptance by Evava Pietri.

[*] Corresponding author at: Department of Psychology, McGill University, 2001 McGill College Ave, Montreal H3A 1G1, Quebec, Canada.
E-mail address: eliane.roy2@mail.mcgill.ca (E. Roy).

https://doi.org/10.1016/j.jesp.2024.104628
Received 21 February 2024; Received in revised form 30 April 2024; Accepted 6 May 2024

Stepanova, Dotsch, & Todorov, 2015), or 4) completed judgments without any additional decision-relevant information (e.g., emotion categorization in Bijlstra, Holland, Dotsch, Hugenberg, & Wigboldus, 2014). A related question then concerns whether people have the capacity to effectively avoid or override this process by limiting reliance on facial information when more judgment-relevant information is available. More specifically, it remains unclear whether individuals can "turn off" or disengage the encoding process to lessen the influence of biasing information on their judgments when given the opportunity to do so.

To that end, the present work explored whether reliance on quickly-extracted social information in judgment persists when participants are 1) never explicitly told about how targets differ on a perceptual trait, 2) placed in a context where most people report not wanting to use such social information when making judgments, and 3) given a decision-making task where more diagnostic information is available. In these more naturalistic contexts, participants may work to effectively override the influence of social information, such as by minimizing the attention given to the biasing cue, which could in turn reduce discriminatory judgment. Similarly, experimentally restricting exposure time to biasing social cues may lessen the cues' influence on judgment. Removing biasing social information may focus attention on decision-relevant cues, which in turn allows people to engage in the more effortful process of impression updating. Indeed, prior work has consistently shown that people can update their initial, automatic impressions when given more diagnostic information (Cone & Ferguson, 2015; Mende-Siedlecki, Baron, & Todorov, 2013).

Thus, limiting exposure to irrelevant social information and directing attention towards more decision-relevant information could then also give people the motivation or ability to update their early impressions based on the more diagnostic qualification information. This hypothesis is also consistent Wilson and Brekke's (1994) mental contamination model, which argues that the contaminating influence of irrelevant information in the evaluation process occurs quickly and effortlessly and that decontamination requires individuals motivation and ability to counteract the influence of these biasing forces.

At the same time, prior work suggests that people may not be able to easily disengage from incorporating social information into their judgment; in one example (Jaeger, Todorov, Evans, & van Beest, 2020), participants made hypothetical sentencing decisions in a context where target faces were hidden and only more diagnostic criteria were presented (e.g., severity of the crime). When participants were then shown an image of the defendant and given the opportunity to update their judgments, those that decided to update their judgment did so in a way that revealed a reliance on facial stereotypes by systematically giving harsher sentences and more judgments of guilt towards faces that were manipulated to appear as more untrustworthy. However, it is unclear whether such effects are removed or weakened among decision-makers who either choose to or are made to spend less time viewing the potentially biasing social information.

To investigate this question, we used an academic decision-making task known to reliably create biases favoring certain social groups over others (e.g., more versus less physically attractive people; Axt & Lai, 2019), despite most participants reporting a desire to avoid using social information – like physical attractiveness – when making decisions (Axt, Nguyen, & Nosek, 2018). Using both correlational and experimental approaches, we then explored whether discrimination is lessened when the attention to – or availability of – biasing information is minimized.

1. Understanding discrimination: sensitivity vs. bias

The present work extends prior applications of the concepts of decision sensitivity and bias, terms rooted in Signal Detection Theory (SDT; Green & Swets, 1966), to discrimination (Axt & Lai, 2019). In judgment contexts that allow for objectively correct or incorrect decisions, sensitivity pertains to the *number* of errors participants make while trying to distinguish between underlying populations (e.g., more versus less qualified applicants). In SDT, sensitivity is indexed by d', where greater values indicates fewer errors. Response bias refers to the degree to which one response is more likely than another, such as whether a hiring manager is more lenient or stringent towards a group. In SDT, response bias is indexed by criterion (c), with lower values indicating more leniency. In the context of social judgment, discrimination can arise when social groups differ in sensitivity (Hinzman, Lloyd, & Maddox, 2022) or response bias (e.g., Correll, Park, Judd, & Wittenbrink, 2002). For example, in one prior study (Axt et al., 2018), participants selecting members of an academic honor society had lower response criterion (i.e., greater leniency) towards applicants from their own university as opposed to another university.

Sensitivity and bias are conceptually and empirically distinct, and the two outcomes may depend on different psychological processes. Various interventions have been shown to differentially impact sensitivity versus bias. In one example, placing time pressure on judgments decreased overall sensitivity but had no effect on criterion biases, (Axt & Lai, 2019). Conversely, warnings about the potential for socially-biased favoritism reduced criterion biases but had no impact on sensitivity (Axt, Casola, & Nosek, 2019).

We hypothesize that sensitivity and bias reflect different aspects of discriminatory judgment (e.g., Lynn & Barrett, 2014; Witt, Taylor, Sugovic, & Wixted, 2015). Specifically, sensitivity is more related to participants' ability or motivation to process decision-relevant information, while bias is more related to participants' ability or motivation to counteract the influence of irrelevant, biasing social information. In many judgment contexts, it is difficult to synthesize relevant judgment criteria (e.g., GPA, interview performance) and arrive at a summary evaluation. Accurately parsing such information requires skill and/or effort, and participants with greater motivation or capacity to process such criteria should commit fewer errors (i.e., higher sensitivity). Consistent with this perspective, Axt and Lai (2019) found that requiring participants to make slower decisions heightened sensitivity, though this work cannot disentangle whether such effects were due to greater processing of decision-relevant information (e.g., applicants' credentials) versus less reliance on decision-irrelevant information (e.g., applicants' faces), since both forms of information were presented for the duration of the judgment.

Conversely, bias should be more related to one's ability or motivation to counteract irrelevant social information. However, unlike the effort needed to parse decision-relevant information, most social information can be gleaned easily, even automatically (Ito & Urland, 2003), and removing biasing social information may then leave behavior unchanged so long as that information has already been encoded. Still, it is possible that reducing exposure to biasing information could be an effective means for altering judgment. For one, a longer presentation of biasing information –such as allowing participants more time to attend to applicants' physical attractiveness – could magnify its impact in decision-making. Indeed, one prior study applied Diffusion Decision Modeling (DDM) to a judgment task that showed favoritism towards physically attractive people (Axt & Johnson, 2021). Here, the DDM drift rate parameter differed between more and less attractive targets, suggesting that the impact of attractiveness on judgment accumulated as decisions unfolded. Limiting participants' ability to view such biasing information may in turn disrupt its' influence on decision-making. This approach could also redirect participant attention towards more decision-relevant information, which would allow participants to update their initial impressions of applicants and thus increase accuracy on the task (i.e., increase sensitivity; Axt & Lai, 2019).

2. Cue availability, bias, and sensitivity

The present work explores these ideas first with an eye-tracking study that investigates the relationship between visual attention and bias and sensitivity in a hypothetical admissions task. Gaze behavior can index the degree to which participants attend to decision-relevant or decision-irrelevant information (Wedel, Pieters, & van der Lans, 2022), factors that often predict judgment (Krajbich, Lu, Camerer, & Rangel, 2012). To date, there have been limited applications of eye-tracking to discrimination research. In one study (Madera & Hebl, 2011), participants watched an interview of an applicant who was facially stigmatized (i.e., had a noticeable scar). Analyses revealed that more time spent looking at the applicant's scar predicted lower ratings of hireability. Buijsrogge, Duyck, and Derous (2021) found similar evidence, as participants gave more initial fixations to fictional job applicants with a facial birthmark.

While this prior work provides insight into the relationship between visual attention and discrimination, the studies used outcome measures that lacked objectively correct answers (i.e., whether applicants were in fact qualified for the job), meaning analyses could not distinguish between the relative contribution of sensitivity and bias. Moreover, data were only correlational, making it ambiguous whether experimentally restricting visual attention through limiting cue availability can impact judgment behavior. It is then unclear whether experimentally restricting visual attention towards social information (i.e., a face) and/or towards decision-relevant information (i.e., qualifications) can impact judgment. To explore this issue, we ran one correlational, eye-tracking study and four experiments that manipulated the extent to which participants could view decision-irrelevant social information versus decision-relevant information.

3. Study 1

3.1. Methods

Participants. 202 undergraduates were recruited from a university pool and received course credit. For this first study, we sought to collect as much data as possible over the course of two semesters. As in prior work using the Judgment Bias Task (JBT; Axt et al., 2018), participants were excluded from analyses if they 1) had an acceptance rate below 20% or above 80%, or 2) accepted or rejected each more or less physically attractive applicant. All studies used the same exclusion criteria to remove participants who either paid insufficient attention to instructions to accept approximately 50% of applicants or those who ignored qualifications and simply accepted or rejected each more or less physically attractive applicant. For all primary analyses across all studies, re-running analyses without any data exclusions changed only one of 49 conclusions (see online supplement for full reporting).

In all, 158 participants in Study 1 had eligible data (88.8% female, 55.1% White, $M_{\text{age}} = 20.5$, $SD_{\text{age}} = 2.20$; see online supplement for full demographics). Forty-three participants were removed due to errors in the eye-tracker calibration,[1] and one for JBT criteria. This sample provided 80% power to detect an effect as small as $r = 0.22$. For all studies, we report all measures and data exclusion. See https://osf.io/rg2st/?view_only=43368f06986a451184d945dc55570b09 for data and materials.

Procedure. In a fixed order, participants completed the JBT while having their gaze recorded with an eye-tracker. In the JBT, participants reviewed 64 applicants for membership into an academic honor society. Each application displayed four pieces of qualification information: 1) Science GPA, 2) Humanities GPA, 3) Interview score, 4) Letter of recommendation quality. Profiles were created such that half were objectively less qualified and half were more qualified. Specifically, each academic qualification could be placed on a 1–4 scale: all objectively more qualified applicants had credentials summing to 14 and all objectively less qualified applicants had credentials summing to 13 (see Axt et al., 2018 for more information on JBT scoring).

Profiles were paired with faces that were pretested to differ in physical attractiveness (Axt et al., 2018), but equated on age and race. Within both levels of applicant quality, 16 profiles were randomly paired with more physically attractive faces (eight men, eight women), and 16 profiles were randomly paired with less physically attractive faces (eight men, eight women). Participants first completed an encoding phase, where they viewed all 64 profiles individually (qualifications only) for one second each. Then, participants completed the decision phase where they were instructed to accept approximately half of the applicants. Participants were given 10 s for each decision before the trial advanced. For each participant, the photos paired with each application were randomly assigned such that applications from each qualification level were matched with 16 (8 men, 8 women) more or less attractive faces.

To ensure the JBT would be adequate for eye-tracking, faces and qualifications took up the same amount of space on the screen. In addition, the vertical placement of the face and qualifications was pseudorandomized across trials, such that half of the trials presented faces in the upper-half of the screen and qualifications in the lower-half, and other half of trials had the reverse orientation. The face stimuli used for more versus less physically attractive faces did not reliably differ in luminance or contrast (see online supplement for more information). Participants were seated at a 24-in. monitor with a resolution of 1280 × 1024 pixels and were instructed to keep their heads rested on a mount positioned 60 cm from the monitor. Participants had the gaze location of the left eye measured using an EyeLink 1000 eye tracker (SR Research, Osgoode, ON) with a sampling rate of 1000 Hz. Stimuli were presented using Inquisit Lab (version 6), which was synchronized with the eye-tracker.

We computed saccades and fixations from raw eye-tracking data using the *pyedfread* package (version 0.1; https://github.com/nwilming/pyedfread) for Python (version 3.8.8), employing default parameters for saccade identification (i.e., minimum velocity threshold of $30^\circ/\text{s}$ and acceleration threshold: $8000^\circ/\text{s}^2$ maintained for at least 4 ms). Movements between two saccades were labelled fixations.

Analyses focused on fixation duration – an indicator of attention and informational processing (Salvucci & Goldberg, 2000)– in two equally-sized areas of interest. Specifically, eye-tracking analyses used the average time participants fixated on either decision-relevant information (the qualifications) or irrelevant information (the face) across trials. See Fig. 1 for sample eye-tracking data. SDT analyses used the scoring procedure outlined in Correll et al. (2002). See online supplement for more information on scoring.

After completing the JBT, participants completed a survey measuring their performance on the task and attractiveness-related attitudes (see online supplement for item wordings and descriptive statistics for all studies). Finally, participants completed a five-item demographics questionnaire that included information on race, age, gender identity, birth country, and first language spoken.

3.2. Results

Task performance. For all participants, we computed the overall JBT acceptance rate and overall JBT accuracy – the percentage of trials where participants correctly accepted more qualified applicants or correctly rejected less qualified applicants. Acceptance rate was nearly 50% of applicants ($M = 51.5\%$, $SD = 9.2\%$), meaning that most participants complied with the task's instructions. Accuracy was above chance ($M = 68.1\%$, $SD = 9.5\%$), and sensitivity (d') was above zero ($M = 1.00$, $SD = 0.55$). Since our distribution of JBT applicants contains an equal number of more and less qualified applicants, accuracy and sensitivity were highly correlated ($r > 0.96$ across all studies). Finally, participants showed a criterion bias, with more attractive applicants ($M = -0.07$, $SD = 0.38$) receiving a lower criterion compared to less attractive applicants ($M = 0.02$, $SD = 0.34$; $t(157) = 2.72$, $p = .007$, $d =$

---

[1] Participants were not pre-screened for having astigmatism or other visual impairments, leading many participants to fail eye-tracker calibration.

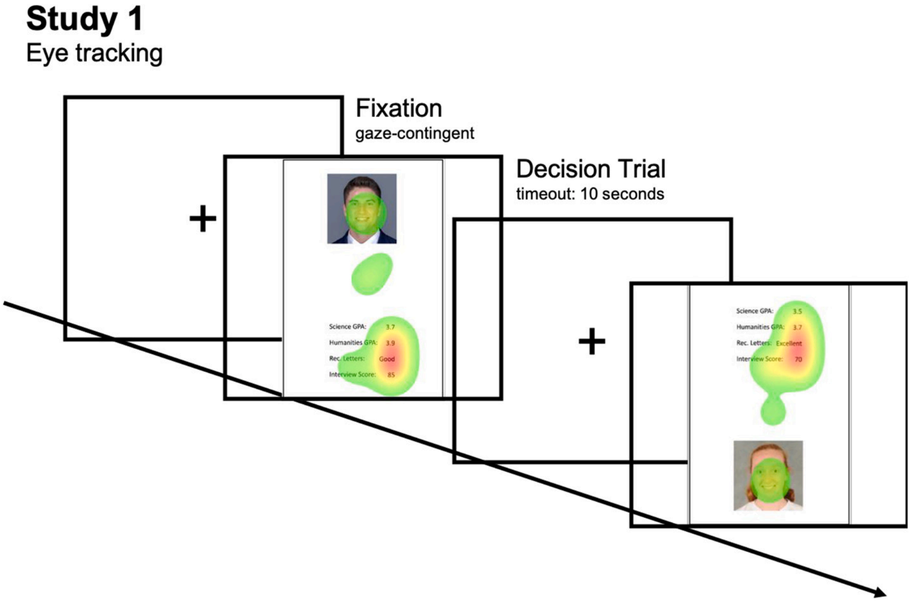

Fig. 1. Study 1 timing and sample eye-tracking data.

0.22, 95% CI [0.06, 0.37]). In other words, participants showed greater leniency towards more versus less physically attractive applicants, despite 87% of participants reported not wanting to use physical attractiveness in their decisions. These findings replicate past work which found similar sensitivity and criterion bias effect in different populations (e.g., Axt et al., 2018).

Eye-tracking analyses. On average, participants spent 2520 ms (*SD* = 1734) per trial to make their decisions. During this time, participants spent significantly less time looking at faces on average ($M = 487$ ms, $SD = 454$) than at qualifications ($M = 2019$ ms, $SD = 1284$; $t(157) = -26.04$, $p < .001$).

We next completed correlational analyses between eye-tracker metrics and JBT overall sensitivity and criterion bias. We used overall sensitivity (sometimes referred to as discriminability; Hyett, Parker, & Breakspear, 2014) given that prior work using the JBT has failed to produce consistent differences in sensitivity between the social groups used in the task (e.g., Axt et al., 2018.). Indeed, in the present work, there were no reliable differences in sensitivity for more versus less physically attractive applicants (all *d*'s < 0.14; see online supplement for more details). As a result, analyses focused on overall sensitivity for ease of interpretation. Overall sensitivity was calculated as the difference between the standardized proportion of *hits* (i.e., the correct acceptances of more qualified applicants) and *false alarms* (i.e., the incorrect acceptances of less qualified candidates). In this sense, a sensitivity value higher than 0 indicates more *hits* relative to *false alarms*, or, in simpler terms, a better ability to correctly accept qualified candidates and correctly reject less qualified applicants.

The criterion value for each subgroup followed prior uses of SDT in similar tasks (e.g., Correll et al., 2002) and was calculated as the sum of standardized *hits* and *false alarms* and then multiplied by −0.5 (see online supplement for the full formula). Here, lower criterion values indicate a lower threshold to give an "accept" response, and positive values indicate a higher threshold. In the context of a distribution where half of the applicants are more versus less qualified, a criterion of zero would represent the neutral bias point where exactly 50% of applicants are accepted. However, since the JBT did not require participants to accept exactly 50% of applicants, the meaning of the zero point is more difficult to interpret, as participants may naturally be more lenient or stringent on average (see Study 1a in Axt & Lai, 2019). A more informative metric is then the relative difference in criterion between more and less physically attractive applicants, which we calculated as the difference between participants' response criterion towards more versus less attractive applicants, with higher values showing greater leniency towards more versus less physically attractive applicants.

On average, time spent looking at qualifications across trials was positively associated with sensitivity ($r(156) = 0.30$, $p < .001$). However, there was no reliable association between average time spent looking at applicant faces and criterion bias ($r(156) = 0.10$, $p = .190$). See Fig. 2 for scatterplots. Moreover, average time looking at qualifications was not related to criterion bias ($r(156) = -0.05$, $p = .522$) and average time looking at faces was not related to sensitivity ($r(156) = 0.10$, $p = .167$).

3.3. Discussion

Study 1 examined different patterns of visual attention in a judgment task that has shown to consistently elicit attractiveness-based discrimination. While time spent looking at decision-relevant information (academic qualifications) was associated with greater sensitivity (i.e., fewer judgment errors), there was no reliable association between time spent looking at faces and bias in criterion (i.e., leniency). That is, the degree to which participants were more lenient towards physically attractive applicants was unrelated to time spent looking at attractiveness-related information.

These findings from Study 1 are consistent with an account that sensitivity in decision-making is associated with greater attention towards – and therefore greater processing of – decision-relevant information. In comparison, bias seems unrelated to the degree of attention paid to social information. We can speculate that since such information is so rapidly encoded (Ito & Urland, 2003; Zhou et al., 2020), there are not further biasing effects on behavior once the facial information has been processed. However, the correlational nature of Study 1 prevents us from making any causal inferences. For one, participants with greater motivation to parse applicant qualifications may be able to do so quite quickly but still choose to attend to such information for a longer time. Conversely, participants with a greater ability to limit the influence of social information may choose to attend to the biasing cue (i.e., faces) since they know that they can override its impact. Stronger evidence would come from experiments manipulating presentation time of social information or decision-relevant information. This was the goal of the remaining studies, which manipulated presentation duration of faces (Studies 2a-2b) or qualifications (Studies 3a-3b).

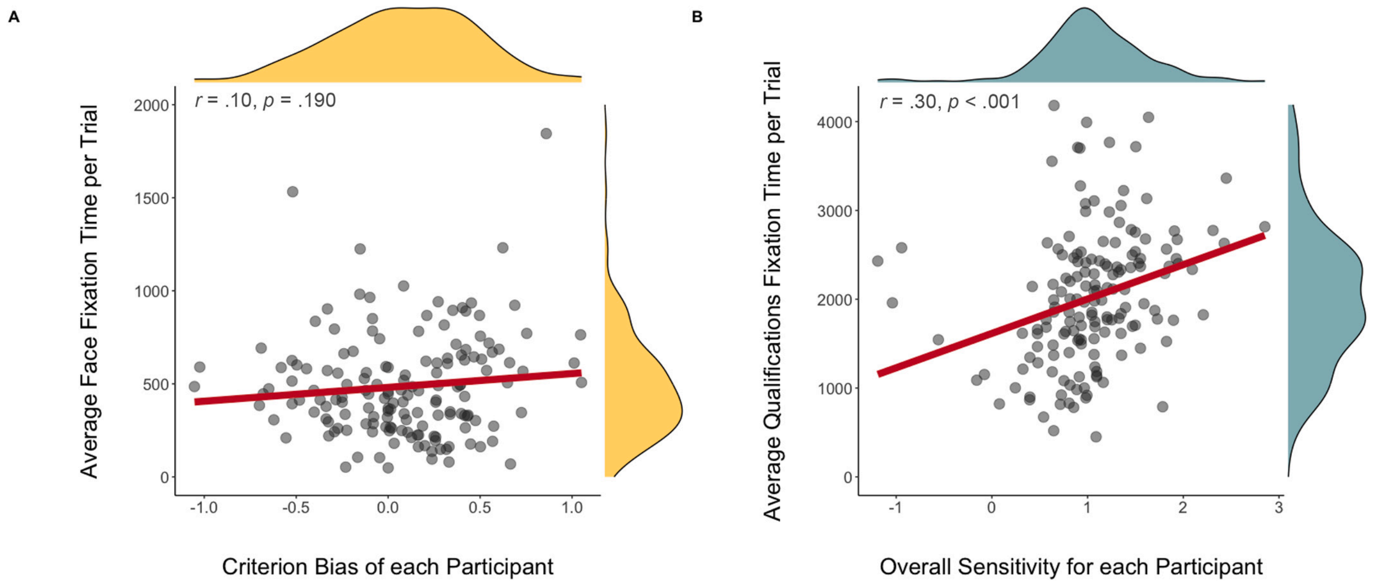

Fig. 2. Relationship between criterion bias, overall sensitivity, and relevant gaze behaviors (Absolute Time) in Study 1.

4. Studies 2a-2b

4.1. Methods

Participants. In Study 2a, 569 participants were recruited from Prolific (UK and USA participants) and received 1.5£ following the study completion. Using the same JBT exclusion criteria from Study 1, we excluded 32 participants based on their JBT performance and one for failing an attention check. At last, 536 participants composed out final sample (47.4% female, 72.6% White, $M_{age} = 37.4$, $SD = 13.7$; see online supplement for full demographics), and this sample provided 80% power to detect an effect as small as $\eta^2 = 0.018$ ($r = 0.13$). In Study 2b, 1035 Prolific participants received 1.5£ following study completion (again restricted to UK and USA participants). Using the same exclusion criteria, the final sample was N = 978 (44.6% female, 84.4% White, $M_{age} = 39.8$, $SD = 13.3$; see online supplement for full demographics), whereas 22 participants were excluded for JBT performance, and two for failing an attention check. This sample provided 80% power to detect an effect as small as $\eta^2 = 0.010$ ($r = 0.099$). In both Studies 2a and 2b, we had set out to achieve 125 eligible participants per condition, but oversampled to be conservative about potential exclusion rates.

Procedure. Participants completed an online, modified version of the JBT described in Study 1 with no eye-tracking. In Studies 2a-2b, applicants' faces disappeared after various amounts of time, while qualifications remained on screen for the full trial duration. In Study 2a, participants were randomly assigned to either the no time constraints condition, where faces were visible for the full trial, or to experimental conditions where faces disappeared after either 500 ms, 2000 ms, or 3500 ms. In Study 2b, participants were randomly assigned to the no time constraints condition or one of two experimental conditions: a condition with a presentation duration so short that faces were unlikely to be encoded (50 ms), or a condition with a presentation duration above previously established thresholds for face encoding (350 ms; Jacques & Rossion, 2006). Participants could decide to accept or reject an applicant at any point during the trial. Participants were assigned to one of 12 JBT orders. Across orders, each face was equally likely to be paired with a more or less qualified profile, and each profile was equally likely to be paired with a more or less physically attractive face.

Lastly, participants completed the same performance survey as Study 1 and a seven-item demographics questionnaire.

4.2. Results

Overall, participants accepted around half of the applicants, as per task instructions (Study 2a: M = 50.7%, SD = 11.3; Study 2b: M = 51.7%, SD = 11.5). They also had above chance accuracy (see Table 1). We next conducted paired-samples t-tests in each condition to determine whether there was a criterion bias between the more versus less physically attractive candidates (see Table 2). In Study 2a, all conditions showed lower criterion for more versus less physically attractive applicants (all d's > 0.30). The same was true of Study 2b, except for the 50 ms condition, where no criterion bias emerged ($d = 0.08$; $p = .314$). See Fig. 3 for plots of JBT performance across all conditions.

We next conducted one-way ANOVAs on JBT sensitivity and criterion bias. Criterion bias referred to the same difference score used in Study 1, such that higher values indicated a greater pro-attractiveness bias. In Study 2a, we found no main effect of condition on sensitivity, $F(3, 532) = 1.42$, $p = .236$, $\eta^2 = 0.008$, or on criterion bias, $F(3, 532) =$

Table 1
Means and standard deviations for JBT metrics in Studies 2a and 2b.

Condition	Accuracy	Sensitivity	More attractive criterion	Less attractive criterion
Study 2a				
No time constraints (N = 138)	0.68 (0.09)	1.03 (0.55)	- 0.06 (0.47)	0.06 (0.43)
500 ms (N = 121)	0.66 (0.09)	0.91 (0.58)	−0.12 (0.43)	0.01 (0.42)
2000 ms (N = 135)	0.67 (0.08)	0.92 (0.50)	−0.12 (0.38)	0.08 (0.42)
3500 ms (N = 142)	0.67 (0.08)	0.98 (0.53)	−0.09 (0.39)	0.07 (0.39)
Study 2b				
No time constraints (N = 334)	0.68 (0.08)	1.02 (0.48)	- 0.10 (0.43)	0.02 (0.41)
50 ms (N = 319)	0.67 (0.08)	1.00 (0.52)	−0.04 (0.40)	−0.02 (0.42)
350 ms (N = 325)	0.67 (0.09)	0.96 (0.54)	−0.14 (0.33)	−0.01 (0.43)

Note. JBT = Judgment Bias Task. More attractive criterion reflects the criterion value applied towards more physically attractive applicants. Less attractive criterion reflects the criterion values applied towards less physically attractive applicants.

Table 2
Within-subjects *t*-tests for criterion bias in each condition for Studies 2a and 2b.

Condition	Criterion Bias	*d*	95% CI
Study 2a			
No time constraints ($N = 138$)	$t(137) = 3.62, p < .001$	0.30	[0.14, 0.48]
500 ms ($N = 121$)	$t(120) = 3.94, p < .001$	0.36	[0.17, 0.54]
2000 ms ($N = 135$)	$t(134) = 4.83, p < .001$	0.42	[0.24, 0.59]
3500 ms ($N = 142$)	$t(141) = 5.2, p < .001$	0.44	[0.26, 0.61]
Study 2b			
No time constraints ($N = 334$)	$t(333) = 5.02, p < .001$	0.27	[0.17, 0.38]
50 ms ($N = 319$)	$t(318) = 0.95, p = .34$	0.05	[−0.06, 0.16]
350 ms ($N = 325$)	$t(324) = 5.00, p < .001$	0.28	[0.17, 0.39]

Note. CI = confidence interval. Criterion bias reflect the difference score between the criterion value for more versus less physically attractive applicants.

0.76, $p = .518$, $\eta^2 = 0.004$. Study 2b found no main effect of condition on sensitivity, $F(2, 975) = 0.98$, $p = .377$, $\eta^2 = 0.002$, but a reliable main effect of condition on criterion bias, $F(2, 975) = 6.37$, $p = .002$, $\eta^2 = 0.013$. To follow up on this, we conducted a series of post-hoc *t*-tests and found that the 50 ms condition had reliably lower criterion bias than the no time constraints condition, $t(651) = 3.15$, $p = .002$, $d = 0.25$, 95% CI [0.10, 0.40], and the 350 ms condition, $t(642) = 3.22$, $p = .001$, $d = 0.25$, 95% CI [0.10, 0.40]. There was no reliable difference between the no time constraints (control) and 350 ms condition, $t(657) = 0.12$, $p = .907$, $d = 0.01$, 95% CI [−0.14, 0.16].

4.3. Discussion

Studies 2a-2b revealed that the duration of exposure to applicant faces did not impact sensitivity and only impacted criterion bias when face presentation was so short (50 ms) as to prevent encoding (e.g., Albert, Wells, Arnocky, Liu, & Hodges-Simeon, 2021; Klatt et al., 2016; Willis & Todorov, 2006). Presenting faces for as little as 350 milliseconds still produced a criterion bias, and one that did not differ in magnitude from a condition where faces were displayed throughout judgment. As such, it seems that quickly removing biasing social information did not provide participants with any greater capacity to lessen the influence of physical attractiveness on decision-making, nor did it redirect attention towards decision-relevant information (i.e., qualifications) in a way that productively increased sensitivity and reduced judgment errors. This finding is interesting especially when considering that for most of the deliberation time during the judgment was spent only looking at the qualification information. Nevertheless, faces were still able to influence the decision-making process.

Following these findings, we wanted to test whether manipulating the amount of time qualification information could be viewed would impact discrimination. Two final studies then take the reverse approach of Studies 2a-2b and investigate how the degree of exposure to decision-relevant information (i.e., qualifications) impacts sensitivity and criterion bias.

5. Studies 3a-3b

5.1. Methods

Participants. In Study 3a, we set out to recruit 125 eligible participants per condition, and oversampled to buffer against any potential data exclusion. We recruited 750 participants from Prolific and compensated them 1.5£ for completion of the study. We used the same exclusion criteria as in Studies 1, 2a, and 2b, and excluded 22 participants based on JBT performance and two for failing the attention check. The final sample of 726 participants (57.4% female, 82.2% White, M_{age}

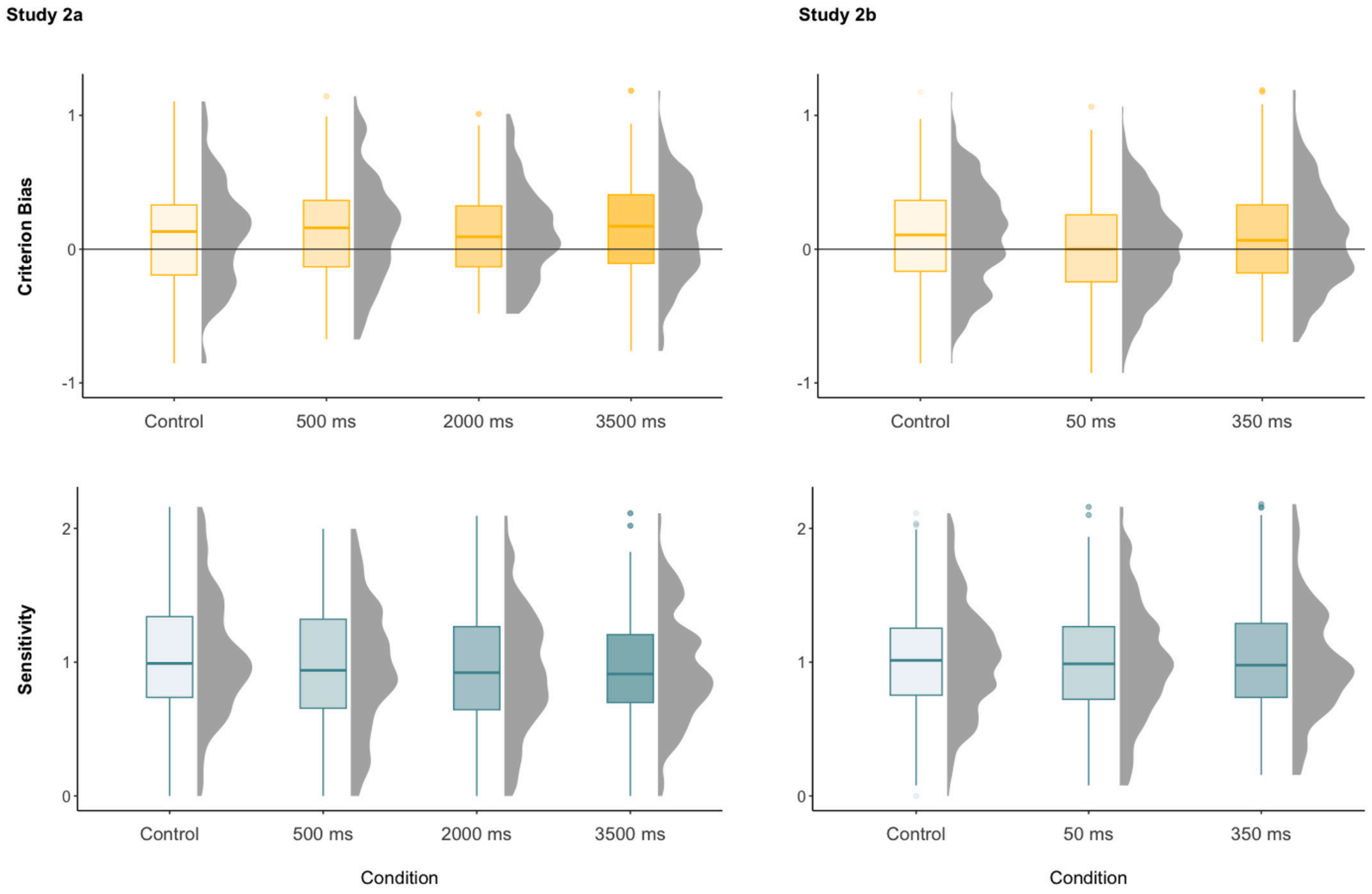

Fig. 3. Mean differences in criterion bias and sensitivity across conditions in Study 2a and 2b.

= 37.9, *SD* = 13.0; see online supplement for full demographics)provided 80% power to detect an effect as small as $\eta^2 = 0.015$ ($r = 0.12$). In Study 3b, 1650 undergraduate participants were recruited from a university pool in exchange for course credit. For this study, we sought to collect as much data as possible over the course of one semester. Using the same data cleaning procedures outlined above, we were left with a sample of 1522 participants (74.1% female, 72.3% White, $M_{age} = 19.7$, SD = 2.9; see online supplement for full demographics); 68 participants were excluded for JBT performance, and 60 for failing an attention check. This sample provided 80% power to detect an effect as small as $\eta^2 = 0.007$ ($r = 0.08$).

Procedure. Studies 3a-3b followed the same overall procedure, in that Study 3b was a replication of Study 3a using a different sample source. First, participants completed a modified version of the JBT in which we manipulated the amount of time qualifications were displayed on the screen. In the experimental conditions, qualifications disappeared after a certain amount of time had elapsed, but faces remained onscreen for the full duration of the trial. In both studies, participants were randomly assigned to either the no time constraints condition, where the qualifications remained visible for the full trial, or to an experimental conditions, where qualifications disappeared after 1500 ms, 3000 ms, or 4500 ms. Following the JBT, participants then answered the same self-report and demographics items as Studies 2a-2b.

5.2. Results

Overall, participants accepted around half of the applicants, as per task instructions (Study 3a: *M* = 53.3%, *SD* = 12.9; Study 3b: *M* = 50.5%, *SD* = 11.6). They also above chance accuracy (see Table 3). We conducted paired-samples *t*-tests in each condition to determine whether there was a criterion bias between the more versus less physically attractive candidates (see Table 4). In Studies 3a and 3b, all conditions showed lower criterion for more versus less physically attractive applicants (all *d*'s > 0.19). See Fig. 4 for plots of JBT performance across all conditions.

We next conducted one-way ANOVAs on JBT sensitivity and criterion bias difference scores. In Study 3a, there was no main effect of condition on criterion bias, $F(3, 722) = 1.26$, $p = .289$, $\eta^2 = 0.005$, but there was a main effect of condition on sensitivity, $F(3, 722) = 5.09$, $p = .002$, $\eta^2 = 0.021$. Follow-up t-tests found that the 1500 ms condition had reliably lower sensitivity compared to the no time constraints condition, $t(354) = 3.73$, $p < .001$ $d = 0.40$, 95% CI [0.19, 0.61], and to the 4500 ms condition, $t(380) = 3.02$, $p = .003$, $d = 0.31$, 95% CI [0.11, 0.51]. No other comparisons were reliable at $p < .05$ (see online supplement for full reporting). In Study 3b, there was no main effect of condition on criterion bias, $F(3, 1518) = 0.45$, $p = .719$, $\eta^2 = 0.001$, but there was a main effect of condition on sensitivity, $F(3, 1518) = 4.57$, $p = .003$, $\eta^2 = 0.009$. Follow-up t-tests showed that the 1500 ms condition had lower sensitivity compared to the no time constraints condition, $t(785) = 3.38$, $p = .001$ $d = 0.24$, 95% CI [0.10, 0.38] and to the 4500 ms condition, $t(740) = 2.57$, $p = .010$, $d = 0.19$, 95% CI [0.05, 0.33]. The 3000 ms condition also had lower sensitivity than the no time constraints condition, $t(778) = 2.42$, $p = .016$ $d = 0.17$. All other comparisons were not significant at $p < .05$ (see online supplement).

Table 3
Means and standard deviations for JBT metrics for Studies 3a and 3b.

Condition	Accuracy	Sensitivity	More attractive criterion	Less attractive criterion
Study 3a				
No time constraints (N = 171)	0.69 (0.08)	1.12 (0.52)	−0.11 (0.40)	0.05 (0.43)
1500 ms (N = 185)	0.66 (0.08)	0.92 (0.47)	−0.14 (0.40)	−0.06 (0.40)
3000 ms (N = 173)	0.68 (0.08)	1.01 (0.51)	−0.09 (0.43)	0.03 (0.44)
4500 ms (N = 197)	0.69 (0.08)	1.08 (0.51)	−0.16 (0.43)	−0.06 (0.46)
Study 3b				
No time constraints (N = 413)	0.68 (0.08)	1.02 (0.47)	−0.08 (0.41)	0.08 (0.41)
1500 ms (N = 374)	0.66 (0.07)	0.91 (0.44)	−0.09 (0.40)	0.07 (0.39)
3000 ms (N = 367)	0.67 (0.08)	0.94 (0.50)	−0.10 (0.40)	0.03 (0.44)
4500 ms (N = 368)	0.67 (0.08)	1.00 (0.48)	−0.08 (0.42)	0.07 (0.44)

Note. JBT = Judgment Bias Task. More attractive criterion reflects the criterion value for more physically attractive applicants. Less attractive criterion reflects the criterion values for less physically attractive applicants.

Table 4
Within-subjects t-tests for criterion bias in each condition for Studies 3a and 3b.

Condition	Criterion Bias	*d*	95% CI
Study 3a			
No time constraints (N = 171)	$t(170) = 4.99$, $p < .001$	0.38	[0.23, 0.54]
1500 ms (N = 185)	$t(184) = 2.57$, $p < .05$	0.19	[0.04, 0.33]
3000 ms (N = 173)	$t(172) = 3.80$, $p < .001$	0.29	[0.14, 0.44]
4500 ms (N = 197)	$t(196) = 3.21$, $p < .01$	0.32	[0.12, 0.52]
Study 3b			
No time constraints (N = 413)	$t(412) = 8.66$, $p < .001$	0.42	[0.33, 0.53]
1500 ms (N = 374)	$t(373) = 8.16$, $p < .001$	0.42	[0.32, 0.53]
3000 ms (N = 367)	$t(366) = 6.36$, $p < .001$	0.33	[0.23, 0.44]
4500 ms (N = 368)	$t(367) = 7.48$, $p < .001$	0.39	[0.28, 0.50]

Note. CI = confidence interval. Criterion bias reflect the difference score between the less and the more attractive criterion values.

5.3. Discussion

In Studies 3a-3b, manipulating the exposure to decision-relevant information (i.e., qualifications) significantly impacted sensitivity but not criterion bias. Across both studies, sensitivity was lower in the 1500 ms condition than in the no time constraints or 4500 ms condition. These results suggest that participants had more difficulty distinguishing between the more versus less qualified applicants when qualifications were only processed for a short time. These results are consistent with the correlational findings from Study 1, in which those participants who spent more time looking at the qualifications across trials also had greater sensitivity. Since applicant qualifications were relatively hard to parse, the evidence from Studies 3a-3b suggests that limiting participants' viewing of such information had the effect of lessening participant's ability to distinguish between more and less qualified applicants.

6. General discussion

Five studies explored how different patterns of visual attention and cue availability were related to distinct components of discrimination. In Study 1, eye-tracking analyses found that time spent viewing decision-relevant qualifications was associated with greater sensitivity but time spent viewing faces was unrelated to biases in response criterion. In four follow-up experimental studies, we provided causal evidence that sensitivity is impacted by presentation duration of more complex, decision-relevant information, though these effects were relatively modest (maximum $r = 0.20$ in Study 3a, maximum $r = 0.12$ in Study 3b). However, bias in response criterion was consistently unrelated to presentation duration of social information, so long as presentation time allowed for encoding. It is notable that visual attention to or cue availability of biasing social information like applicant faces did not contribute to the magnitude of discrimination in judgment, suggesting that other avenues are needed in the search for strong predictors of discriminatory behavior (Kurdi et al., 2019).

These studies extend prior research concerning just how rapidly

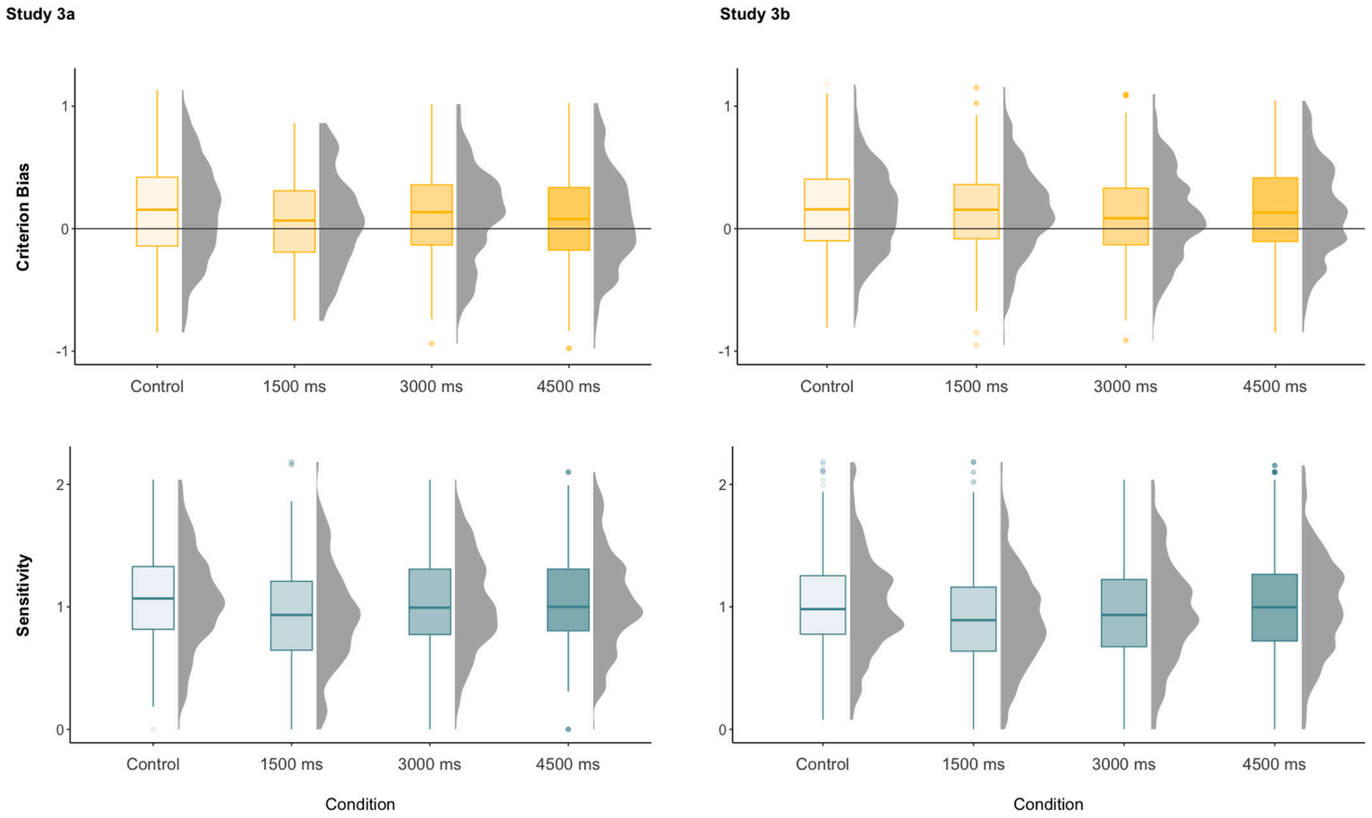

Fig. 4. Mean differences in criterion bias and sensitivity across conditions in Study 3a and 3b.

social information can be extracted from faces. For instance, trait judgments like trustworthiness or competence can be rather accurately extracted from faces in as little as 100 ms (Ambrus, Kaiser, Cichy, & Kovács, 2019; Domen, Derks, Van Veelen, & Scheepers, 2020; Hehman, Sutherland, Flake, & Slepian, 2017; Todorov et al., 2009). Consistent with this work, Studies 2a-2b revealed that attractiveness-based discrimination emerged when faces were presented for as little as 350 milliseconds. Yet, it is striking that, in comparison to the studies cited above, participants in our studies 1) were not instructed to attend to applicant attractiveness, 2) were provided information (i.e., qualifications) that was actually relevant to judgment, 3) had ample opportunity to update their impressions of applicants after faces were removed, and 4) largely reported wanting to avoid using such social information when making their decisions (i.e., over 90% of participants indicated they did not want to use physical attractiveness when completing the JBT). However, discriminatory judgment persisted despite a goal of avoiding physical attractiveness, a very brief exposure to stimuli that communicated attractiveness, and the presence of decision-relevant information throughout judgment. While prior studies have shown that accurate face processing can be "turned on" in very brief windows of time, these data suggest that such face processing also cannot be "turned off", even when participants have a limited window in which to view faces, a desire to avoid using facial information in judgment, and other relevant information to attend to. These findings may be best explained by the ease at which facial cues are processed, rather than an inherent belief that these cues are informative to the judgment at hand (Jaeger, Evans, Stel, & van Beest, 2019).

This work also adds to our understanding of how attractiveness-based discrimination arises in the JBT. For instance, past research (Axt & Lai, 2019, Study 1a) found that completely removing faces from the JBT increased sensitivity relative to a control condition, perhaps because this "blind" condition led to deeper processing of the decision-relevant criteria. This result is consistent with research in visual attention, which finds that simultaneously presenting different types of information (e.g., faces, qualifications) will create attention competition, and removing one source of information leads to better processing of the remaining information (e.g., Ashby, Walasek, & Glöckner, 2015; Teoh, Yao, Cunningham, & Hutcherson, 2020). From this perspective, it is notable that the 50 ms condition in Study 3a showed no differences in sensitivity relative to the no time constraints condition. The 50 ms condition may then not have pushed participants to better parse the academic qualifications; instead, participants might have still attempted to encode applicants' faces but simply could not do so within the short presentation window. As a result, the present work reveals how even presenting social information for a very short amount of time can still elicit attention competition that suppresses judgment accuracy.

In addition, the similarity in JBT performance among conditions presenting applicants' faces for at least 350 milliseconds in Studies 2a-2b sheds light on prior JBT analyses that used diffusion decision modeling (Axt & Johnson, 2021), where drift rate parameters suggested that physical attractiveness exerted an effect on judgment throughout the decision-making process. That removing faces in Studies 2a-2b – often well before judgments were made – failed to change behavior suggests that the influence of social information is not dependent on such information being actively available. Rather, such information must simply have been encoded *at some point* during the decision-making process (Wilson & Brekke, 1994), and the influence of these information on judgment accumulates as decisions unfold. These findings also speak to the many contexts in which physical attractiveness cues impact judgments (Feingold, 1992). One possible explanation for this effect is evolutionary: as postulated in Zebrowitz (2011), attractiveness-based trait inferences appear to always have been a part of social life, and incorporating attractiveness information into the judgments studied here may reflect some downstream consequences of prior evolutionary adaptations for tasks like finding mates with high genetic fitness. That is, while participants completing this form of the JBT have largely indicated a desire to avoid using physical attractiveness when judging applicants (Axt et al., 2018), the continued use of such information may signal reliance on a cue that had some function in our past and as a result is not so easily overridden.

The present work makes clear that merely manipulating exposure to socially biasing information is unlikely to change behavior. From a practical perspective, these results could inform efforts to reduce discrimination, such as by ruling out intuitively appealing but likely unproductive interventions. For instance, the present results would suggest that conducting job interviews by phone or with video cameras off would have little impact if applicants were initially selected for interviews based off LinkedIn profiles that communicated biasing information like age, attractiveness or prestige of undergraduate institution. To that end, these results cast doubt on any evaluation process that allows for judgments to be made after potentially biasing social information is encoded – even if it is later removed – as doing so may fail to reduce discrimination or increase judgment accuracy. In contexts where the complete removal of potentially biasing social information (i.e., “blinding”; Goldin & Rouse, 2000) is impossible or undesired (Doleac & Hansen, 2020), we would expect that merely limiting the presentation of such information to have no consistent impact on discriminatory behavior. Rather, researchers and practitioners may have to look towards interventions that provide decision-makers with the capacity to override the influence of such social information (Morewedge et al., 2015) or that restructure the decision-making environment to lessen the chances that such biases will be allowed to operate (e.g., creating shortlists; Lucas, Berry, Giurge, & Chugh, 2021).

At the same time, these studies show that directing attention towards decision-relevant information is an effective way to reduce errors, and thereby decrease overall discrimination. For example, even in contexts where bias among evaluators is held constant, simulation studies shows that increasing accuracy in decision (e.g., correctly promoting the most deserving employees) can results in meaningful discrimination reduction, especially when effects are compounded over time (Meldgin, Mitchell, & Oswald, 2024). This recommendation to focus on sensitivity-increasing strategies would align with efforts in San Francisco to automatically remove information like race from police reports and only allow prosecutors access to decision-relevant information (e.g., severity of crime; McDede, 2019). In this sense, completely removing such irrelevant information could be expected to not only lessen racial bias but also increase attention given to more judgment-relevant criteria (Axt & Lai, 2019). However, sensitivity-increasing avenues for discrimination are often overlooked; in one prior JBT study (Axt, Yang, & Deshpande, 2023), only 30% of lay participants thought requiring a response delay would improve performance on the task, while 56% (wrongly) believed that reviewing a lesson on the dangers of confirmation bias would be effective.

Nonetheless, our conclusions are limited by use of a single outcome measure, the JBT, and by the features of the task. First, the sole reliance on the JBT as an outcome measure of discrimination means that conclusions may not generalize to other contexts, though separate work (Correll, Wittenbrink, Crawford, & Sadler, 2015) found similar results of bias persisting despite brief cue availability when using a separate outcome tasks that involved racial biases in the recognition of weapons versus harmless objects (though in these studies the availability of decision-relevant and irrelevant information was removed simultaneously).

In addition, since the outcome used in this work was artificial in nature, we cannot infer whether the cue availability has the same impact of sensitivity in real-world decision contexts, where social cues are more multi-faceted than only a picture as presented in the JBT. In addition, future work should explore whether effects emerge in contexts where acceptance rates are more competitive (e.g., graduate student admissions), as a more competitive context could elicit different patterns of behaviors (e.g., Krosch, Tyler, & Amodio, 2017). Follow-up studies will also want to address the lack of pre-registration for these studies, and the exclusive focus on physical attractiveness. Although an impactful domain of discrimination (e.g., Lippens et al., 2023), physical attractiveness is by no means the only domain in which discrimination occurs. Finally, since evaluation processes may be domain-dependent, future work will want to investigate how sensitivity and bias relates to the availability of difference social cues, and in contexts where multiple biases operate simultaneously (e.g., Axt et al., 2019).

Social information – like physical attractiveness – is processed efficiently and effortlessly. Combatting discrimination by limiting exposure to such social cues after it has been encoded reflects an optimistic but inaccurate understanding of social cognition. Effectively reducing social judgment biases will then come from either promoting greater processing of decision-relevant information, or from providing more concrete strategies that allow people to overcome the influence of biasing information on their behavior.

Open practices statement

Materials, data, and analysis syntax for all studies is available at: https://osf.io/rg2st/?view_only=43368f06986a451184d945dc55570b09.

CRediT authorship contribution statement

Eliane Roy: Writing – review & editing, Writing – original draft, Visualization, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Y. Doug Dong:** Writing – review & editing, Formal analysis, Data curation. **A. Ross Otto:** Writing – review & editing, Resources, Methodology, Conceptualization. **Jordan Axt:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

None.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jesp.2024.104628.

References

Albert, G., Wells, E., Arnocky, S., Liu, C. H., & Hodges-Simeon, C. R. (2021). Observers use facial masculinity to make physical dominance assessments following 100-ms exposure. *Aggressive Behavior, 47*(2), 226–235. https://doi.org/10.1002/ab.21941

Ambrus, G. G., Kaiser, D., Cichy, R. M., & Kovács, G. (2019). The neural dynamics of familiar face recognition. *Cerebral Cortex, 29*(11), 4775–4784.

Ashby, N. J. S., Walasek, L., & Glöckner, A. (2015). The effect of consumer ratings and attentional allocation on product valuations. *Judgment and Decision making, 10*(2), 172–184. https://doi.org/10.1017/S1930297500003934

Axt, J. R., Casola, G., & Nosek, B. A. (2019). Reducing social judgment biases may require identifying the potential source of bias. *Personality and Social Psychology Bulletin, 45* (8), 1232–1251. https://doi.org/10.1177/0146167218814003

Axt, J. R., & Johnson, D. J. (2021). Understanding mechanisms behind discrimination using diffusion decision modeling. *Journal of Experimental Social Psychology, 95*, Article 104134. https://doi.org/10.1016/j.jesp.2021.104134

Axt, J. R., & Lai, C. K. (2019). Reducing discrimination: A bias versus noise perspective. *Journal of Personality and Social Psychology, 117*(1), 26–49. https://doi.org/10.1037/pspa0000153

Axt, J. R., Nguyen, H., & Nosek, B. A. (2018). The judgment bias task: A flexible method for assessing individual differences in social judgment biases. *Journal of Experimental Social Psychology, 76*, 337–355. https://doi.org/10.1016/j.jesp.2018.02.011

Axt, J. R., Yang, J., & Deshpande, H. (2023). Misplaced intuitions in interventions to reduce attractiveness-based discrimination. *Personality and Social Psychology Bulletin, 49*, 527–540.

Bijlstra, G., Holland, R. W., Dotsch, R., Hugenberg, K., & Wigboldus, D. H. J. (2014). Stereotype associations and emotion recognition. *Personality and Social Psychology Bulletin, 40*, 567–577.

Bruneau, E., Szekeres, H., Kteily, N., Tropp, L. R., & Kende, A. (2020). Beyond dislike: Blatant dehumanization predicts teacher discrimination. *Group Processes & Intergroup Relations, 23*(4), 560–577.

Buijsrogge, A., Duyck, W., & Derous, E. (2021). Initial impression formation during the job interview: Anchors that drive biased decision-making against stigmatized applicants. *European Journal of Work and Organizational Psychology, 30*(2), 305–318. https://doi.org/10.1080/1359432X.2020.1833980

Cañadas, E., Rodríguez-Bailón, R., Milliken, B., & Lupiáñez, J. (2013). Social categories as a context for the allocation of attentional control. *Journal of Experimental Psychology: General, 142*(3), 934.

Chen, J. M., & Hamilton, D. L. (2012). Natural ambiguities: Racial categorization of multiracial individuals. *Journal of Experimental Social Psychology, 48*(1), 152–164.

Chwe, J. A., & Freeman, J. B. (2023). Trustworthiness of crowds is gleaned in half a second. *Social Psychological and Personality Science*. https://doi.org/10.1177/19485506231164703

Cone, J., & Ferguson, M. J. (2015). He did what? The role of diagnosticity in revising implicit evaluations. *Journal of Personality and Social Psychology, 108*(1), 37.

Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology, 83*(6), 1314–1329. https://doi.org/10.1037/0022-3514.83.6.1314

Correll, J., Wittenbrink, B., Crawford, M. T., & Sadler, M. S. (2015). Stereotypic vision: How stereotypes disambiguate visual stimuli. *Journal of Personality and Social Psychology, 108*(2), 219.

Doleac, J. L., & Hansen, B. (2020). The unintended consequences of "ban the box": Statistical discrimination and employment outcomes when criminal histories are hidden. *Journal of Labor Economics, 38*(2), 321–374.

Domen, I., Derks, B., Van Veelen, R., & Scheepers, D. (2020). Gender identity relevance predicts preferential neural processing of same-gendered faces. *Social Neuroscience, 15*(3), 334–347.

Dunham, Y., Stepanova, E. V., Dotsch, R., & Todorov, A. (2015). The development of race-based perceptual categorization: Skin color dominates early category judgments. *Developmental Science, 18*(3), 469–483.

Feingold, A. (1992). Good-looking people are not what we think. *Psychological Bulletin, 111*(2), 304–341. https://doi.org/10.1037/0033-2909.111.2.304

Forscher, P. S., Lai, C. K., Axt, J. R., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2019). A meta-analysis of procedures to change implicit measures. *Journal of Personality and Social Psychology, 117*(3), 522–559. https://doi.org/10.1037/pspa0000160

Goldin, C., & Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review, 90*(4), 715–741.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics* (Vol. 1, pp. 1969–2012). New York: Wiley.

Hehman, E., Sutherland, C. A. M., Flake, J. K., & Slepian, M. L. (2017). The unique contributions of perceiver and target characteristics in person perception. *Journal of Personality and Social Psychology, 113*, 513–529.

Hinzman, L., Lloyd, E. P., & Maddox, K. B. (2022). The stigmatized perceiver: Exploring the implications of social stigma for cross-race face processing and memory. *Social and Personality Psychology Compass, 16*(2), Article e12654.

Hyett, M., Parker, G., & Breakspear, M. (2014). Bias and discriminability during emotional signal detection in melancholic depression. *BMC Psychiatry, 14*, 1–12.

Ito, T. A., & Urland, G. R. (2003). Race and gender on the brain: Electrocortical measures of attention to the race and gender of multiply categorizable individuals. *Journal of Personality and Social Psychology, 85*(4), 616.

Jacques, C., & Rossion, B. (2006). The speed of individual face categorization. *Psychological Science, 17*(6), 485–492. https://doi.org/10.1111/j.1467-9280.2006.01733.x

Jaeger, B., Evans, A. M., Stel, M., & van Beest, I. (2019). Explaining the persistent influence of facial cues in social decision-making. *Journal of Experimental Psychology: General, 148*(6), 1008–1021. https://doi.org/10.1037/xge0000591

Jaeger, B., Todorov, A. T., Evans, A. M., & van Beest, I. (2020). Can we reduce facial biases? Persistent effects of facial trustworthiness on sentencing decisions. *Journal of Experimental Social Psychology, 90*, Article 104004. https://doi.org/10.1016/j.jesp.2020.104004

Kauff, M., & Wagner, U. (2012). Valuable therefore not threatening: The influence of diversity beliefs on discrimination against immigrants. *Social Psychological and Personality Science, 3*(6), 714–721.

Klatt, T., Maltby, J. J., Humphries, J., Smailes, H. L., Ryder, H., Phelps, M., & Flowe, H. D. (2016). Looking bad: Inferring criminality after 100 ms. *Applied Psychology in Criminal Justice, 12*(2), 114–125. http://www.apcj.org/journal/index.php?mode=view&item=118.

Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology, 3*, 193. https://doi.org/10.3389/fpsyg.2012.00193

Krosch, A. R., Tyler, T. R., & Amodio, D. M. (2017). Race and recession: Effects of economic scarcity on racial discrimination. *Journal of Personality and Social Psychology, 113*(6), 892.

Kurdi, B., Seitchik, A. E., Axt, J. R., Carroll, T. J., Karapetyan, A., Kaushik, N., … Banaji, M. R. (2019). Relationship between the implicit association test and intergroup behavior: A meta-analysis. *American Psychologist, 74*(5), 569–586. https://doi.org/10.1037/amp0000364

Lucas, B. J., Berry, Z., Giurge, L. M., & Chugh, D. (2021). A longer shortlist increases the consideration of female candidates in male-dominant domains. *Nature Human Behaviour, 5*(6), 736–742.

Lynn, S. K., & Barrett, L. F. (2014). "Utilizing" signal detection theory. *Psychological Science, 25*, 1663–1673.

Madera, J., & Hebl, M. (2011). Discrimination against facially stigmatized applicants in interviews: An eye-tracking and face-to-face investigation. *The Journal of Applied Psychology, 97*, 317–330. https://doi.org/10.1037/a0025799

McDede, H. J. (2019, September 9). *New software takes race out of the picture. Does it make prosecution more fair?* KALW. https://www.kalw.org/show/crosscurrents/2019-09-09/new-software-takes-race-out-of-the-picture-does-it-make-prosecution-more-fair.

Meldgin, D. R., Mitchell, G., & Oswald, F. L. (2024). Modeling gender differences in the job promotion process: Replication and extension of Martell et al. (1996). *Journal of Applied Psychology*. https://doi.org/10.1037/apl0001179. Advance online publication.

Mende-Siedlecki, P., Baron, S. G., & Todorov, A. (2013). Diagnostic value underlies asymmetric updating of impressions in the morality and ability domains. *Journal of Neuroscience, 33*(50), 19406–19415.

Milkman, K. L., Akinola, M., & Chugh, D. (2012). Temporal distance and discrimination: An audit study in academia. *Psychological Science, 23*(7), 710–717. https://doi.org/10.1177/0956797611434539

Morewedge, C. K., Yoon, H., Scopelliti, I., Symborski, C. W., Korris, J. H., & Kassam, K. S. (2015). Debiasing decisions: Improved decision making with a single training intervention. *Policy Insights From the Behavioral and Brain Sciences, 2*(1), 129–140. https://doi.org/10.1177/2372732215600886

Munder, A. K., Becker, J. C., & Christ, O. (2020). Standing up for whom? Targets' different goals in the confrontation of discrimination. *European Journal of Social Psychology, 50*(7), 1443–1462.

Rooth, D. O. (2009). Obesity, attractiveness, and differential treatment in hiring: A field experiment. *Journal of Human Resources, 44*, 710–735.

Salvucci, D. D., & Goldberg, J. H. (2000). Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the 2000 symposium on eye tracking research & applications* (pp. 71–78). https://doi.org/10.1145/355017.355028

Teoh, Y. Y., Yao, Z., Cunningham, W. A., & Hutcherson, C. A. (2020). Attentional priorities drive effects of time pressure on altruistic choice. *Nature Communications, 11*(1). https://doi.org/10.1038/s41467-020-17326-x. Article 1.

Todorov, A., Pakrashi, M., & Oosterhof, N. N. (2009). Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition, 27*(6), 813–833. https://doi.org/10.1521/soco.2009.27.6.813

Wedel, M., Pieters, R., & van der Lans, R. (2022). Modeling eye movements during decision making: A review. *Psychometrika*, 1–33.

Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological Science, 17*(7), 592–598.

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

Witt, J. K., Taylor, J. E. T., Sugovic, M., & Wixted, J. T. (2015). Signal detection measures cannot distinguish perceptual biases from response biases. *Perception, 44*(3), 289–300.

Zebrowitz, L. A. (2011). Ecological and social approaches to face perception. In A. J. Calder, et al. (Eds.), *Oxford handbook of face perception* (pp. 31–50). Oxford University Press.

Zhou, Y., Gao, T., Zhang, T., Li, W., Wu, T., Han, X., & Han, S. (2020). Neural dynamics of racial categorization predicts racial bias in face recognition and altruism. *Nature Human Behaviour, 4*(1), 69–87.