

Journal of Personality and Social Psychology

What Counts as Discrimination? How Principles of Merit Shape Fairness of Demographic Decisions

Teodora K. Tomova Shakur and L. Taylor Phillips

Online First Publication, January 13, 2022. <http://dx.doi.org/10.1037/pspi0000383>

CITATION

Tomova Shakur, T. K., & Phillips, L. T. (2022, January 13). What Counts as Discrimination? How Principles of Merit Shape Fairness of Demographic Decisions. *Journal of Personality and Social Psychology*. Advance online publication. <http://dx.doi.org/10.1037/pspi0000383>

What Counts as Discrimination? How Principles of Merit Shape Fairness of Demographic Decisions

Teodora K. Tomova Shakur and L. Taylor Phillips
Stern School of Business, New York University

Demographic attributes (e.g., age, disability, race) frequently affect people's decisions. We provide a novel perspective as to why such discrimination persists: Meritocratic principles lead people to perceive some demographic attributes as fair to use, rather than as discriminatory. Specifically, we theorize that meritocracy requires that controllable and relevant inputs determine outcomes; as a result, perceived controllability and relevance affect the degree to which demographic attributes are perceived as fair to use. Moreover, we suggest perceived relevance *outweighs* controllability, such that even uncontrollable attributes can be perceived as fair criteria if perceived to be outcome relevant. In two qualitative studies, we probed how people think about demographic attributes used in selection (Studies 1a–b). We find that people refer to controllability and relevance dimensions to justify their perceptions. Further, people largely associate uncontrollable, irrelevant attributes with discrimination (race, sex), neglecting attributes they perceive as controllable and/or relevant (disability, caregiving status). Next, three surveys (Studies 2a–c) support our theorizing that perceived relevance impacts fairness perceptions more strongly than perceived controllability. In three experiments (Studies 3a–c), we provide causal evidence that relevance and controllability shape perceived fairness, which in turn affects selection behaviors, including seeking information regarding demographic attributes during hiring. Finally, Study 4 demonstrates downstream consequences: Perceived controllability, relevance, and use of demographic attributes together impact employees' psychological safety and job satisfaction. Overall, we find that principles of merit lead people to believe that even some legally protected demographic attributes are fair to use, allowing discrimination to persist.

Keywords: discrimination, meritocracy, fairness, legitimacy, diversity

Supplemental materials: <https://doi.org/10.1037/pspi0000383.supp>

People make social interaction decisions every day: whom to hire, to whom to sell, whom to promote, and whom to date. And in these daily interactions, discrimination and intergroup inequity can be perpetuated. For instance, hiring, housing, and myriad other social decisions continue to be illegally influenced by protected demographic attributes,¹ such as sex and gender (Crosby, 2008; Eagly & Heilman, 2016; Heilman, 1983), race (Bertrand & Mullainathan,

2004; Kang et al., 2016; Pager & Shepherd, 2008), and age (Abrams et al., 2016; Finkelstein et al., 1995). Such decisions are also influenced by nonprotected demographic attributes, which can nevertheless contribute to inequity. For instance, people are more likely to select candidates from elite educational institutions, even when other candidates have better qualifications (Collins, 2019; Rivera, 2016).

Previous work has focused on the role that explicit and implicit biases may play in such discriminatory decisions (Arrow, 1972; Bertrand & Mullainathan, 2004; Danbold & Huo, 2017; Pager & Shepherd, 2008; Rosette, Akinola, & Ma, 2018; see Wax, 1999 for legal lens). These perspectives suggest that bias can lead people to subvert meritocratic principles, in order to hire the candidates they prefer (e.g., Norton et al., 2004; Uhlmann & Cohen, 2005). Thus, current efforts often aim to increase awareness of and reduce lurking biases as a means to reduce discrimination (e.g., Devine et al., 2012). In contrast, we suggest an additional reason discrimination may

Teodora K. Tomova Shakur  <https://orcid.org/0000-0003-4798-8012>
L. Taylor Phillips  <https://orcid.org/0000-0003-4655-6904>

We thank Elizabeth Campbell, Sora Jun, Lisa Leslie, Sean Malahy, Michael North, Elad Sherf, Batia Wiesenfeld, and the Meritocracy & Inequality Lab for helpful comments on earlier versions of this article. We are grateful for research assistance from Khalid Tomova Shakur and New York University (NYU) Stern SPUR students. We appreciate research support provided by the NYU Stern Center for Global Economy and Business. Study materials are provided in the Online Supplemental Material.

Teodora K. Tomova Shakur played a lead role in investigation, formal analysis, and writing of original draft. L. Taylor Phillips played a lead role in conceptualization, methodology, validation, and writing review and editing; and an equal role in writing of original draft.

Correspondence concerning this article should be addressed to Teodora K. Tomova Shakur and L. Taylor Phillips, Stern School of Business, New York University, 44 West 4th Street, New York, NY 10012, United States. Email: ttomova@stern.nyu.edu and ltaylorphillips@stern.nyu.edu

¹ Here, we use “demographic attributes” to reflect classifiable characteristics extracted from major social categories. We investigate a range of demographic attributes, including frequently studied attributes such as gender and race, alongside understudied others such as political affiliation and religion. Importantly, each of the demographic attributes we probe has been found to affect discrimination: Abrams et al. (2016), Bertrand and Mullainathan (2004), Collins (2019), Eagly and Heilman (2016), Ghumman and Jackson (2010), Gift and Gift (2015), Hoffman (2011), Jones (2013), Krefting and Brief (1976), Lee et al. (2018), Nadler and Kufahl (2014), Pitesa and Pillutla (2019), Rhode (2010), Rivera (2016), and Seidel et al. (2000).

persist is because people find selection based on some kinds of demographic attributes to be meritocratic in the first place.

How do individuals evaluate which demographic attributes are fair to use and which are not? Here, we theorize that meritocratic principles may contribute. In Western contexts especially, strong prescriptive norms toward meritocracy shape perceptions of fairness (Castilla & Benard, 2010; Deutsch, 1975; Kluegel & Smith, 1986; Phillips & Lowery, 2020; Son Hing et al., 2011; Uhlmann & Cohen, 2005, 2007). As such, we argue that meritocratic principles constrain what *kinds* of inputs are appropriate: specifically, principles of controllability and relevance are foundational to meritocracy. In turn, we suggest that these foundational principles will shape which demographic attributes are perceived to be fair to use in selection, rather than discriminatory.

Appraising Discrimination

Existing work has considered when and why individuals recognize objectively discriminatory treatment as such. For instance, the attributions to discrimination framework suggests that people perceive a decision as discriminatory when it is seen as based on a social identity and as unjust (Major, Quinton, & McCoy, 2002; Major & Dover, 2016). In turn, scholars have identified myriad factors that affect whether a decision is perceived to be based on social identity and unjust, including individuals' own group membership (Major et al., 2003), beliefs about the target (Rodin et al., 1990), perceptions of intent (Apfelbaum et al., 2017), and the framing of the decision as focusing on victims or beneficiaries (Phillips & Jun, 2021).

However, Jetten, Iyer, et al. (2013) note that this work on individual reactions to discrimination has largely focused on discriminatory experiences that are both legally and normatively considered unfair or illegitimate² (e.g., race or sex in the present, Western milieu). In contrast, when social identity or demographic selection is perceived to be *legitimate*, those discriminated against are less likely to engage in corrective collective action (Jetten et al., 2011). Moreover, the antecedents to such understandings—how a domain of discrimination comes to be understood as unfair—has received even less attention (Jetten, Iyer, et al., 2013; see also Jetten, Branscombe, et al., 2013). That is, even when a decision is clearly based on social identity, individuals may not *perceive* it to be unjust, and thus may not recognize it as discriminatory.

Here, we take up this call to explore the antecedents of perceived fairness of differential treatment on the basis of demographic attributes. In Western contexts, the logic of meritocracy undergirds individuals' understanding of legitimacy and fairness (Davey et al., 1999; Son Hing et al., 2011). As such, we begin with meritocracy and suggest that this logic will shape when and why people perceive specific demographic attributes to be fair to use, rather than discriminatory.

Requisites of Merit: Controllability and Relevance

The term meritocracy was coined to envision a social system in which highly desired outcomes (e.g., prosperous jobs) are obtained through individual talent and effort; although the word was intended as satire (Young, 1958), meritocracy has come to represent the core of the American dream (Hochschild, 1996; Kluegel & Smith, 1986; Son Hing et al., 2011). Indeed, research has found high levels of belief in and support for meritocracy in the U.S.: People believe that

individuals' positive attributes (e.g., hard work) will be proportionally rewarded, while negative attributes (e.g., laziness) will be punished accordingly, which in turn offers a sense of predictability and justice (Son Hing et al., 2011; see also Adams, 1965; Deutsch, 1975; Lerner, 1977, 1980).

How does one know if meritocracy is being upheld? Meritocracy enshrines the equity principle of distributive justice: Meritocracy requires inputs to be proportional to outcomes received (Adams, 1965; Walster et al., 1973). However, we emphasize here a neglected additional view: Meritocracy also specifies what *kinds* of inputs should be considered. For example, individually controllable inputs are critical to many definitions of meritocracy, across the social sciences. McCoy and Major (2007, p. 341) describe meritocracy as “locating the responsibility for social status within the efforts and abilities of individuals,” and McNamee and Miller (2009) define meritocracy as a system in which “individuals get ahead and earn rewards in direct proportion to their individual efforts and abilities.” Likewise, Davey et al. (1999) emphasize meritocracy as involving internally controllable inputs, especially effort, hard work, and ability.

Moreover, in his original work defining the equity principle as a ratio of inputs to outcomes, Adams (1965, p. 277) describes relevance as a critical characteristic of inputs: “If he perceives it to be relevant, if he expects a just return for it, it is an input.” Likewise, Son Hing et al. (2011, p. 433) write that meritocracy is idealized in part because “only relevant inputs . . . should be considered and irrelevant factors . . . should be ignored when distributing outcomes.” In their review, Belmi et al. (2022) suggest that relevance in particular is the most neglected principle of meritocracy.

In short, meritocracy prescribes not only a distributive system but also specific procedural principles: Inputs must be controllable and relevant (Adams, 1965; Belmi et al., 2022; Son Hing et al., 2011). As such, we suggest that people use these two principles to determine whether an attribute counts as merit and in turn to decide whether selection based on that attribute is fair or discriminatory. Moreover, we suggest that relevance will *outweigh* controllability, leading relevant but uncontrollable attributes to be seen as fair criteria. Below, we draw on both justice and attribution literatures to further elaborate these core principles of merit.

Controllability

Controllability, also referred to as “internality” or “achieved” characteristics (Foladare, 1969; Kelley, 1967; Ross, 1977), suggests that attributes must be earned by the individual through their own efforts and volition. This is as opposed to ascribed characteristics, or “accidents of birth” (Foladare, 1969, p. 53), which encompass many demographic attributes. For instance, both justice and attribution literatures suggest that when outcomes are attributable to internal qualities over which individuals have direct control (e.g., winning a

² Fairness has been theorized as the perception of the global appropriateness and legitimacy of an outcome (Colquitt & Rodell, 2015). This conceptualization of fairness is also consistent with how lay people define fairness, including references to justice, legitimacy, appropriateness, morality, and rule adherence (Lupfer et al., 2000). Relatedly, legitimacy has been defined as “the extent to which an outcome is perceived as fair and just” (Bylsma et al., 1995, p. 223; see also Major, 1994; Tyler, 2006). We follow this precedent and treat fairness as a global indicator of appropriateness, justice, and legitimacy.

race with a fast-running time rather than speedy shoes), then people perceive such outcomes as meritorious and deserved (Feather, 1999, 2006; Kelley, 1967; Phillips & Lowery, 2020; Ross, 1977; Thibaut & Walker, 1975). From an attributional standpoint, controllability signals this personal responsibility for outcomes (Feather, 1999; Weiner, 1995). Indeed, controllability may be so valorized as fundamental to meritocracy because internal control helps satisfy fundamental needs for autonomy and predictability (e.g., Adams, 1965; Beauvois & Dubois, 1988; Kay et al., 2009; Thibaut & Walker, 1975; van den Bos & Lind, 2002).

Attributions to discrimination work itself has focused on controllability of a demographic identity in particular as an important predictor of perceptions of stigma, essentialism, and discrimination (Haslam et al., 2000; Hegarty & Golden, 2008; Major & Dover, 2016; Major, Quinton, & McCoy, 2002; Marti et al., 2000; Pachankis et al., 2018; Rothbart & Taylor, 1992). For instance, when people believe that obesity is a controllable attribute, they feel that stigma against obese individuals is deserved, fair, and acceptable (Hegarty & Golden, 2008; Major, 1994; Quinn & Crocker, 1999). Relatedly, Jetten, Iyer, et al. (2013) present a theoretical model of potential antecedents to perceiving discrimination as legitimate, including permeability of group boundaries. That is, to the extent group membership or status can be changed or controlled by the individual, then decisions based on that membership may be seen as more tolerable. Other work finds that believing an attribute to be biological and therefore uncontrollable increases essentialism (Haslam et al., 2000; Prentice & Miller, 2007). While debate persists about whether controllability beliefs increase stigma, or vice versa (Crandall, 1994; Crandall & Eshleman, 2003; Haslam & Levy, 2006; Hegarty & Golden, 2008; see also Salter et al., 2016), we suggest that reliance on meritocratic principles itself should result in individuals finding more controllable attributes to be more fair to select upon.

In addition, controllability is foundational to meritocracy. As such, we expect people will use perceived controllability of a demographic attribute to determine whether an attribute reflects merit and is therefore fair to use in selection.

Relevance

Moving beyond controllability, our meritocratic principles lens further suggests that relevance is another key dimension for determining whether an attribute is fair to use in selection. *Relevance*, also referred to as “job relatedness” and “face validity” (Arvey & Sackett, 1993; Gilliland, 1993; Smither et al., 1993), suggests that attributes must be specifically related to performance outcomes, communicating information about expected performance (Adams, 1965; Gilliland, 1995; Son Hing et al., 2011; see also Berger & Fişek, 2006). For instance, someone who is applying for a data scientist position should be judged based on their analytical skills, rather than their party planning experience (despite that fact that both are individually controllable).

The justice literature has focused on the importance of relevance for procedural fairness (Gilliland, 1993, 1995; Smither et al., 1993). Specifically, job tests and interview criteria are considered relevant when they relate to individuals’ capabilities to perform the job and perform it well. However, whereas this previous work has focused especially on tests designed specifically to evaluate job skills (e.g., work sample tests, Gilliland, 1995), we suggest perceivers will

apply similar relevance standards to the use of demographic attributes, rather than merely perceiving these as categorically irrelevant. Indeed, given many applicants are first excluded based on demographic attributes (e.g., sexual orientation, Dipboye & Colella, 2005) before they even reach the stage of skills tests, people’s assessments of the relevance of such demographic criteria are important to consider.

Interestingly, the attributions to discrimination literature has not generally considered relevance. However, supporting our contentions, Jetten, Iyer, et al. (2013) offer an exception, suggesting in a theoretical model that the perception that exclusion criteria and group membership are related may increase perceptions of legitimacy. That is, to the extent that a particular performance exclusion criteria (physical strength) is perceived to be relevant to group membership (sex), then even members of the disadvantaged group may find the attribute legitimate for selection.

In addition, relevance is also foundational to meritocracy. As such, we expect that, alongside perceived controllability, people will also use perceived relevance of a demographic attribute to determine whether an attribute reflects merit and is therefore fair to use in selection.

Dealing With Ambiguity: Relevance Outweighs Control

We have highlighted classic sociological and psychological perspectives, which describe how meritocracy requires that both controllable and relevant inputs determine outcomes. Following this logic, we suggest that when people perceive a demographic attribute as highly controllable and highly relevant, they will perceive selection based on that attribute as fair. However, some demographic attributes may be perceived as relevant but uncontrollable or irrelevant but controllable. For example, religious beliefs may be perceived as irrelevant for most jobs, but simultaneously be viewed as relatively controllable. Alternatively, caregiving status may be perceived as relevant to some jobs (e.g., likelihood of working overtime), but also perceived as uncontrollable. How might such mismatch between relevance and controllability ultimately affect individuals’ perceptions of which attributes are perceived as fair to use in decision-making?

We suggest that people will rely more on perceived relevance of the demographic attribute than perceived controllability when assessing fairness, in part because selection contexts—for instance, hiring, promotion, leasing, and more—may evoke a business-case mindset. Specifically, previous work shows that a business case mindset can lead people to prioritize material costs and benefits over moral reasoning (Belmi & Pfeffer, 2015; Belmi & Schroeder, 2021; Miller, 2001; see also Noon, 2007). For instance, when people are asked to use a “business case” perspective, they tend to justify and even engage in more unethical behavior (Wang et al., 2014; Zhong, 2011). This mindset becomes intensified in many selection environments because such environments direct people’s attention to material benefits, especially maximizing gains and minimizing losses (Belmi & Pfeffer, 2015; Belmi & Schroeder, 2021; Kouchaki et al., 2013).

Moreover, although this relevance dimension has been largely neglected in work on meritocracy beliefs (as reviewed above), there is reason to consider it the primary principle involved in moving away from other forms of “ocracy” (organization by wealth, organization by family line) toward organization by performance

ability. To the extent relevance then is the first principle of meritocratic organization, it may carry more weight. That is, given meritocratic selection tasks are particularly concerned with fair prediction (see also Lassetter et al., 2021; Prentice & Miller, 2007), then relevance may win out over control. Thus, we suggest that people may perceive relevant demographic attributes as fair to use during selection, even if they are perceived as uncontrollable, because *relevance* prioritizes performance costs and benefits. Following this logic, we suggest people will weigh perceived relevance more heavily than perceived controllability when judging the fairness of using demographic attributes.

Merit Principles May Justify Consequences

Given their commitment to meritocracy, people should be motivated either to engage in meritocratic practices or to justify their practices as meritocratic (e.g., Castilla & Benard, 2010; Son Hing et al., 2011; Uhlmann & Cohen, 2005, 2007). As such, meritocracy serves as a critical legitimizing ideology, through which people can launder selfish preferences, biased motivations, or systemic injustice. For instance, commitment to meritocratic ideologies can lead individuals to deny they benefit from unfair privileges (Knowles & Lowery, 2012; Phillips & Lowery, 2020) and to deny that group differences in society are unjust (Ledgerwood et al., 2011; McCoy & Major, 2007; Savani & Rattan, 2012; Son Hing et al., 2011). As another example, when given a set of criteria (e.g., language fluency vs. math ability), individuals are more likely to rate the skills of their preferred candidate as reflecting the most merit, relative to other skills (Uhlmann & Cohen, 2005; see also Dunning et al., 1989; Norton et al., 2004).

While this previous work focuses on the role of meritocracy to justify preexisting preferred outcomes, we suggest that principles of meritocracy might also serve an antecedent function: To the extent meritocratic principles shape people's belief in the fairness of using demographic attributes, it should also shape their willingness to act. Specifically, perceiving a demographic attribute as controllable and relevant, and therefore fair, should increase people's likelihood of *using* that attribute in selection decisions. We test this possibility in three ways.

First, to the extent that people find an attribute fair to use in selection, we argue that individuals will express more curiosity about job applicants' status regarding those attributes. Relatedly, we expect that individuals will act on this curiosity and actively *seek* information about such attributes, driven by the belief that they are fair to use. That is, when demographic attributes are perceived to be controllable and relevant, people will be more curious and willing to seek out information regarding those attributes in others.

Second, to the extent that people find an attribute fair to use in selection, then they should be more willing to take said attribute into consideration when making decisions, thus affecting people's actual *evaluation* of candidates who possess such attributes. If people find a demographic attribute unfair to use, then they should be more likely to ignore this information during the evaluation process, even when the information itself is present. Indeed, individuals can modulate their attention to irrelevant information to some degree, tuning out information they deem impertinent (McNamara & McDaniel, 2004). Further, individuals should be more likely to base their decisions on criteria they find fair, in part because such decisions are more easily justified (see Norton et al., 2004; Uhlmann &

Cohen, 2005). Thus, when demographic attributes are perceived to be controllable and relevant, those attributes should exert a bigger effect on evaluations.

Third, to the extent that people find an attribute fair to use in selection, people may be more (vs. less) willing to support organizations which use said attributes. That is, perceived fairness of organizational selection practices should have important downstream consequences: When people experience environments they perceive to be unfair, their psychological experience should suffer. Unfair environments are less predictable and decrease feelings of autonomy and control (van den Bos & Lind, 2002). Indeed, procedural justice scholars emphasize that fairness of selection processes can lead to a positive candidate experience, even when the selection outcome is ultimately a rejection (Brockner et al., 1994; Gilliland, 1995). Likewise, discrimination scholars have found that noticing discriminatory practices, even against others in the organization, depresses employees' sense of psychological safety, satisfaction, and commitment (Goldman et al., 2006; Purdie-Vaughns et al., 2008; Sanchez & Brock, 1996; Sanchez et al., 2017; Schmitt et al., 2014). In short, perceiving discrimination and unfairness can be stressful, even when the observer is not the direct target. As a result, working in an environment that is perceived to be unfair should depress employee outcomes. Following this logic, we expect that perceptions of controllability and relevance of demographic attributes used in selection in one's *own* organization will affect psychological safety and job satisfaction regarding that organization.

The Current Research

We have hypothesized that two core principles of meritocracy—controllability and relevance—will shape people's perceptions of the fairness of using demographic attributes in selection decisions. Further, we expect perceived *relevance* of demographic attributes will exert a stronger effect on fairness perceptions than will perceived controllability. In turn, we suggest that these perceptions of fairness will affect people's willingness to discriminate, by leading people to seek out and even use demographic information in selection decisions. Finally, we expect such perceptions of demographic attributes will shape important outcomes (psychological safety and job satisfaction) when they observe their organizations using said attributes.

We explore these ideas in nine studies, using both qualitative and quantitative methods. In Studies 1a and 1b (preregistered), we consider which attributes employees perceive to be fair versus discriminatory in selection processes, and why. Building on these results, in Studies 2a, 2b (preregistered), and 2c, we test perceived controllability and relevance as key mechanisms driving fairness perceptions. In Studies 3a (preregistered), 3b, and 3c, we then experimentally manipulate perceptions of the controllability and relevance of an attribute and test effects on discriminatory behavior: willingness to use the attribute in selection. Finally, in Study 4, we survey employed adults and assess how perceptions of controllability and relevance of demographic attributes used in their own organizations may relate to downstream consequences: employees' psychological safety and job satisfaction.

We aim to make two key contributions. First, we highlight a neglected view of meritocracy: Beyond proportionality, principles of controllability and relevance dictate what *kinds* of inputs are

deemed meritocratic. We use this view to account for perceptions of fairness across a range of demographic attributes. In doing so, we consider a wide variety of well-studied (e.g., race) and under-studied (e.g., religion, caregiving status; see also Iyer et al., 2014) demographic attributes which affect decision-making, ultimately identifying which attributes are seen as more versus less discriminatory. Our approach helps integrate separate work on attributions to discrimination and justice perceptions, by highlighting how both controllability and relevance work in tandem to shape impressions across a range of attributes.

Second, we help shed light on why discrimination may persist. Focusing on core principles of merit, we show that many legally protected attributes are perceived to be controllable and/or relevant, and that relevance especially affects perceptions of fairness. As a result, such attributes are more likely to be used in selection. Indeed, perceiving a demographic attribute as controllable or relevant can make people more tolerant of discrimination, even when it is otherwise illegal. As such, we demonstrate that discrimination is reinforced even through core philosophies of merit: People may discriminate because they think such actions are *fair*, rather than discriminatory.

Study 1a

As a first step, we used a qualitative, open-ended approach to probe the factors driving lay perceptions of fairness of using demographic attributes. Specifically, we explored what rationales people provided for using (or not using) demographic attributes in selection. We relied on a between-subjects design, offering participants one of 15 possible attributes for reflection.

Method

Participants

We recruited 450 U.S. adults from Prolific to complete the study in exchange for \$.45. We removed incomplete and duplicate responses (based on platform-based unique participant identifier), resulting in a final sample of 444 participants. See Table 1 for demographic information across all studies.³

Procedure

Participants were randomly assigned to read about and rate one of 15 possible demographic attributes: age, alma mater, caregiving responsibilities, disability, educational attainment, family origin (i.e., last name), gender, national origin, personal network connections, physical attractiveness, political affiliation, race, religious beliefs, sexual orientation, or socioeconomic status. They first read a prompt describing the use of the target attribute in selection processes:

Think about [race]-based selection in the workplace. This form of selection involves treating someone unfavorably because they belong to a certain [racial] group, or because of personal attributes associated with [race] (e.g., [their appearance]). For instance, a representative of a particular [race] might not be selected for a job due to their [race].

Then, participants rated the perceived fairness of using the attribute in selection decisions. Participants next were asked to explain their rationale: "Please explain your rationale below.

That is, why do you think that this form of selection is [piped in participant fairness rating]."

We reminded participants of their own fairness rating (e.g., "not at all fair") to focus them on explaining the reasons behind rating the attribute as fair or unfair (rather than restating fairness itself as a rationale). Finally, participants reported demographics.

Measures

Perceived Fairness. We measured perceived fairness on a scale from 1 (*not at all*) to 7 (*extremely*) with the item: "This type of selection is fair."

Fairness Rationale Coding. We measured fairness rationale by content-coding participants' open-ended justifications, specifically exploring whether *controllability* and *relevance* were referenced. Two independent coders (naive to hypotheses) coded the open-ended responses using these two themes, indicating either the presence (coded as 1) or the absence (coded as 0) of a theme. Additionally, in the presence of a theme, the coders indicated whether participants described low controllability/relevance (-1), no direction (0), or high controllability/relevance (1). The two coders demonstrated high interrater reliability (Cohen's $k = .70-1$), and thus we combined their codes to create composites for the presence and for the direction of each theme.

Results

See Table 2 for means and standard deviations across demographic attributes and studies.

Fairness Rationales

First, we considered which themes were present. In line with our theorizing, we found that either the controllability, relevance, or both themes were mentioned in 62.16% of all responses. Also, in line with our theorizing, we found that perceived relevance was mentioned more frequently (54.05%) than controllability (13.32%; see Table 3 for specific examples).

Of the responses that coders rated as not mentioning either theme (37.84%), the vast majority used tautological logic (i.e., explained fairness with fairness; "I don't think it is fair or unfair so I choose moderately fair" or "yeah, I think it's a good deal fair"). Moreover, no additional themes emerged from the independent coders' notes.

Second, we considered the direction of the mentioned themes. We specified a mixed-model regression, assessing the effect of high, neutral, or low controllability rationale (fixed effect) on participants' ratings of perceived fairness, including a random-intercept effect of type of demographic attribute. As expected, referring to an attribute as more controllable was significantly associated with perceived fairness, $b = .76$, $SE = .11$, $t(16) = 7.12$, $p < .001$. We also specified the same mixed-effects model, but assessing the effect of high,

³ Our hypotheses focus on meritocracy concepts common in a U.S. cultural context (Son Hing et al., 2011). Thus, across all studies, we restricted all samples to adults located in the U.S. We used prescreening to ensure that participants met our demographic requirements for each specific study (e.g., employed adults in Study 4). Further, we removed duplicate participants and incomplete observations before analysis. Across studies, we aimed for 100 per cell in experimental designs, following current best practice recommendations.

Table 1
Participant Demographics: Studies 1a–4

Variable	Study 1b					Study 2c	Study 2b	Study 3a	Study 3b	Study 3c	Study 4
	Study 1a	Phase 1	Phase 2	Study 2a	Study 2b						
Race											
Asian	10.81%	13.83%	12.29%	7.13%	4.07%	14.85%	4.07%	11.33%	11.08%	6.60%	7.67%
Black	8.11%	4.26%	4.66%	6.52%	8.24%	6.52%	8.24%	6.32%	8.03%	9.20%	7.67%
Latino	7.43%	5.32%	5.87%	5.78%	0.05%	6.44%	0.05%	6.75%	4.99%	6.60%	3.00%
Other	5.63%	6.38%	7.54%	6.65%	3.51%	3.86%	3.51%	9.15%	6.09%	7.78%	6.33%
White	68.02%	70.21%	69.37%	63.22%	78.79%	68.33%	78.79%	66.23%	69.53%	69.58%	75.33%
Decline to state	—	—	0.28%	10.70%	0.09%	—	0.09%	0.22%	0.28%	0.24%	—
Gender											
Female	51.35%	67.02%	51.12%	51.91%	48.67%	46.44%	48.67%	55.56%	52.08%	49.76%	54.33%
Male	46.40%	31.91%	47.21%	37.15%	50.28%	52.02%	50.28%	42.27%	46.54%	48.82%	45.67%
Nonbinary	2.25%	1.06%	1.58%	.25%	1.04%	1.55%	1.04%	2.18%	1.39%	1.42%	—
Decline to state	—	—	.09%	10.70%	—	—	—	—	—	—	—
Age	34.70 (13.00)	33.86 (12.86)	35.06 (13.23)	36.49 (11.95)	36.72 (11.59)	33.24 (12.81)	36.72 (11.59)	31.29 (11.13)	33.09 (11.95)	29.06 (9.60)	34.62 (11.53)

neutral, or low relevance rationale. As expected, we found that referring to an attribute as more relevant was significantly associated with perceived fairness, $b = .70$, $SE = .08$, $t(164) = 8.72$, $p < .001$.

Discussion

In Study 1a, we examined participants' open-ended justifications regarding what renders a selection decision based on a demographic attribute fair or unfair. As expected, participants' rationales suggest that perceived controllability and relevance play an important role; indeed, relevance was mentioned in the majority of responses, and controllability also emerged frequently. Notably, we find a similar pattern in an inductive study, using a response-driven thematic extraction approach to generate a list of potential rationale themes (see Online Supplemental Material). Finally, connecting this qualitative data to participants' quantitative fairness ratings, we also find initial evidence that individuals' perceptions of a demographic attribute as controllable and relevant are related to perceiving selection based on that attribute as fair.

Study 1b

Next, we used a converse logic approach to probe how people think about the use of demographic attributes in selection. That is, Study 1a tested whether, when given different demographic attributes, people's fairness rationales vary in terms of controllability and relevance. In contrast, Study 1b tests whether, when given a fair versus unfair scenario, people *generate* demographic attributes that vary in terms of controllability and relevance.

We used a yoked design. In the first phase of the study, participants generated lists of attributes that they perceived as influencing either a hiring or a discrimination incident. In the second phase, a new set of participants (naive to first phase) rated the controllability and relevance of these attributes. We expected that participants in the first phase would generate controllable, relevant attributes (as rated by second-phase participants) in response to a hiring incident, and uncontrollable, irrelevant attributes in response to a discrimination incident. Study 1b was preregistered: <https://aspredicted.org/y7jj8.pdf>.

Method

Phase 1

Participants. In the first phase of the study, we recruited 100 new U.S. adults from Prolific to complete the study in exchange for \$.75. We removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), resulting in 94 participants in total.

Procedure. We used a within-subjects, two condition design. First, participants were instructed to imagine that a hiring manager at a fictitious company was interviewing two job candidates for an accounting position. Then, participants were introduced to the Hiring condition: They were asked to list all attributes that might affect the manager's decision to hire either candidate in an open-ended, essay format. Afterward, participants were asked to list attributes that might affect the decision using a different question format (15 single blank slots) to encourage a thorough and thoughtful attribute generation.

Table 2*Means and Standard Deviations (Studies 1a, 2a, 2b, 2c, and 4)*

Attribute	Perceived controllability	Perceived relevance	Perceived fairness
Past work performance	S4: 5.73 (1.41)	S4: 5.82 (1.29)	S4: 5.89 (1.30)
Work experience	S4: 4.75 (1.66)	S4: 5.43 (1.39)	S4: 5.52 (1.43)
Educational attainment	S2a: 4.28 (1.52)	S2a: 5.04 (1.50)	S1a: 4.37 (1.25)
	S2b: 3.46 (1.20)	S2b: 3.66 (1.28)	S2a: 5.06 (1.47)
	S2c: 3.52 (1.05)	S2c: 3.9 (1.34)	S2b: 4.43 (1.19)
	S4: 5.41 (1.40)	S4: 5.26 (1.35)	S2c: 4.65 (1.3)
Networks	S4: 5.30 (1.58)	S4: 5.30 (1.58)	S4: 5.30 (1.58)
	S2a: 3.95 (1.53)	S2a: 2.63 (1.56)	S1a: 2.36 (1.22)
	S2b: 3.69 (1.55)	S2b: 3.01 (1.59)	S2a: 3.41 (1.39)
	S2c: 3.76 (1.48)	S2c: 2.91 (1.49)	S2b: 3.20 (1.44)
	S4: 5.16 (1.44)	S4: 3.52 (1.67)	S2c: 3.34 (1.40)
Alma mater	S4: 5.16 (1.44)	S4: 3.52 (1.67)	S4: 2.81 (1.66)
	S2a: 2.65 (1.50)	S2a: 2.66 (1.61)	S1a: 1.87 (1.33)
	S2b: 2.62 (1.21)	S2b: 2.42 (1.44)	S2a: 3.13 (1.41)
	S2c: 2.86 (1.22)	S2c: 2.58 (1.55)	S2b: 3.12 (1.47)
Age*	S2c: 2.86 (1.22)	S2c: 2.58 (1.55)	S2c: 2.83 (1.46)
	S2a: 1.55 (1.31)	S2a: 2.68 (1.49)	S1a: 2.31 (1.26)
	S2b: 1.40 (.85)	S2b: 2.80 (1.31)	S2a: 2.85 (1.42)
	S2c: 1.42 (1.08)	S2c: 2.77 (1.45)	S2b: 3.08 (1.30)
	S4: 1.52 (1.40)	S4: 2.86 (1.62)	S2c: 3.02 (1.46)
Caregiving*	S4: 1.52 (1.40)	S4: 2.86 (1.62)	S4: 2.08 (1.47)
	S2a: 2.65 (1.57)	S2a: 2.31 (1.40)	S1a: 3.59 (1.68)
	S2b: 3.17 (1.46)	S2b: 3.38 (1.43)	S2a: 2.82 (1.44)
	S2c: 3.18 (1.43)	S2c: 3.59 (1.80)	S2b: 3.81 (1.61)
Disability*	S2c: 3.18 (1.43)	S2c: 3.59 (1.80)	S2c: 3.74 (1.85)
	S2a: 1.57 (1.20)	S2a: 2.88 (1.52)	S1a: 2.35 (1.23)
	S2b: 1.59 (.91)	S2b: 2.85 (1.40)	S2a: 2.53 (1.39)
	S2c: 1.49 (.93)	S2c: 2.43 (1.13)	S2b: 2.92 (1.57)
	S4: 1.61 (1.28)	S4: 2.98 (1.67)	S2c: 2.62 (1.59)
Family origin	S4: 1.61 (1.28)	S4: 2.98 (1.67)	S4: 1.99 (1.47)
	S2a: 1.65 (1.24)	S2a: 1.78 (1.29)	S1a: 1.32 (1.02)
	S2b: 1.62 (1.28)	S2b: 1.66 (1.21)	S2a: 2.59 (1.34)
	S2c: 1.86 (1.40)	S2c: 1.78 (1.30)	S2b: 2.16 (1.27)
Physical attractiveness	S2c: 1.86 (1.40)	S2c: 1.78 (1.30)	S2c: 2.16 (1.24)
	S2a: 3.02 (1.42)	S2a: 2.03 (1.34)	S1a: 1.36 (.73)
	S2b: 2.82 (1.49)	S2b: 2.06 (1.22)	S2a: 2.50 (1.30)
	S2c: 2.77 (1.23)	S2c: 1.82 (1.05)	S2b: 2.28 (1.20)
Political affiliation	S2c: 2.77 (1.23)	S2c: 1.82 (1.05)	S2c: 2.30 (1.32)
	S2a: 3.88 (1.54)	S2a: 1.83 (1.32)	S1a: 1.43 (.68)
	S2b: 4.25 (1.14)	S2b: 2.03 (1.27)	S2a: 2.57 (1.29)
	S2c: 4.09 (1.23)	S2c: 2.73 (1.52)	S2b: 2.24 (1.15)
Socioeconomic status	S2c: 4.09 (1.23)	S2c: 2.73 (1.52)	S2c: 3.31 (1.67)
	S2a: 2.96 (1.35)	S2a: 2.20 (1.34)	S1a: 1.92 (1.47)
	S2b: 2.55 (1.27)	S2b: 1.94 (1.18)	S2a: 2.11 (1.27)
	S2c: 2.94 (1.31)	S2c: 2.33 (1.42)	S2b: 2.08 (1.21)
National origin*	S2c: 2.94 (1.31)	S2c: 2.33 (1.42)	S2c: 2.17 (1.23)
	S2a: 1.58 (1.27)	S2a: 1.75 (1.28)	S1a: 1.52 (1.15)
	S2b: 1.43 (1.07)	S2b: 1.52 (1.06)	S2a: 2.38 (1.47)
	S2c: 1.44 (.86)	S2c: 1.53 (1.04)	S2b: 2.29 (1.40)
Religious beliefs*	S2c: 1.44 (.86)	S2c: 1.53 (1.04)	S2c: 1.96 (1.11)
	S2a: 3.26 (1.61)	S2a: 1.65 (1.25)	S1a: 1.38 (.82)
	S2b: 3.38 (1.52)	S2b: 1.62 (.92)	S2a: 2.19 (1.30)
	S2c: 3.41 (1.09)	S2c: 1.65 (.96)	S2b: 2.23 (1.33)
Gender*	S2c: 3.41 (1.09)	S2c: 1.65 (.96)	S2c: 2.13 (1.30)
	S2a: 1.58 (1.22)	S2a: 1.91 (1.37)	S1a: 2.27 (1.68)
	S2b: 1.94 (1.25)	S2b: 1.66 (1.13)	S2a: 2.18 (1.33)
	S2c: 2.30 (1.46)	S2c: 1.77 (1.22)	S2b: 2.15 (1.24)
	S4: 1.69 (1.48)	S4: 1.98 (1.51)	S2c: 2.16 (1.44)
Sexual orientation [†]	S4: 1.69 (1.48)	S4: 1.98 (1.51)	S4: 1.62 (1.40)
	S2a: 2.09 (1.63)	S2a: 1.67 (1.32)	S1a: 1.04 (.19)
	S2b: 1.75 (1.29)	S2b: 1.35 (.88)	S2a: 2.06 (1.31)
	S2c: 2.00 (1.53)	S2c: 1.51 (1.17)	S2b: 1.78 (1.14)
	S4: 2.35 (2.00)	S4: 1.50 (1.16)	S2c: 1.87 (1.50)
			S4: 1.55 (1.37)

(table continues)

Table 2 (continued)

Attribute	Perceived controllability	Perceived relevance	Perceived fairness
Race*	S2a: 1.47 (1.16)	S2a: 1.60 (1.28)	S1a: 1.33 (.83)
	S2b: 1.40 (1.07)	S2b: 1.42 (1.04)	S2a: 1.83 (1.23)
	S2c: 1.62 (1.41)	S2c: 1.86 (1.51)	S2b: 1.93 (1.23)
	S4: 1.50 (1.43)	S4: 1.68 (1.46)	S2c: 1.98 (1.42)
			S4: 1.51 (1.30)

Note. The asterisk (*) indicates legally protected attributes in the U.S. The cross symbol (†) indicates attributes with debated legal protection at the time of some data collection.

Next, participants were introduced to the Discrimination condition: They read information that one of the job candidates believed they had been discriminated against, despite both candidates having similar strengths and weaknesses. Participants were then instructed to generate a list of all forms of discrimination that could have occurred in an open-ended, essay format. Afterward, they also responded using 15 single blank slots to encourage naming as many discrimination types as above. Finally, participants completed demographic information.

Phase 2

Participants. In the second phase of the study, we recruited 1,500 new U.S. adults from Prolific to complete the study in exchange for \$.30. Following our preregistration, we removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), and individuals who failed a simple antibody check, resulting in a final sample of 1,074 participants.

Procedure. Participants were given a list of 10 attributes (randomly selected from 738 unique attributes generated in Phase 1).⁴ For each attribute, participants rated their perceived controllability and perceived relevance.

Measures

Independent Variable. *Hiring versus Discrimination condition* was manipulated in Phase 1, within-subjects, as described above.

Dependent Variables. *Perceived controllability* ($r = .58$) was measured in Phase 2, using a composite of two items: “To what extent [is the target] characteristic” . . . “under someone’s personal control” and “easily changeable” (see Hegarty & Golden, 2008; 1 = *not at all*, 7 = *extremely*).

Perceived relevance ($r = .63$) was measured in Phase 2, using a composite of two items: “To what extent [is the target] characteristic” . . . “relevant to most jobs” and “a signal of someone’s capabilities” (see Bauer et al., 2001; 1 = *not at all*, 7 = *extremely*).

Results

See Table 4 for correlations across studies.

In the Hiring condition, individuals generated $M = 12.93$ attributes. The most frequently mentioned attributes included job candidates’ work experience ($n = 51$), education ($n = 41$), personality ($n = 38$), age ($n = 29$), and attitude ($n = 20$; see Table S1, Online Supplemental Material). The least frequently mentioned attributes included examples like clothing ($n = 1$) and wearing glasses ($n = 1$).

In the Discrimination condition, individuals generated $M = 7.09$ attributes. The most frequently mentioned attributes included race ($n = 117$), age ($n = 86$), gender ($n = 85$), religion ($n = 37$), and disability ($n = 31$). The least frequently mentioned attributes included examples like job candidates’ abilities ($n = 1$), incompetence ($n = 1$), and inexperience ($n = 1$).

Next, we specified a mixed-model regression, assessing the effects of Phase 1 condition (Hiring = 1, Discrimination = -1) on Phase 2 perceived controllability, including a random-intercept of item-generating (Phase 1) participant. As expected, attributes generated in the Hiring condition were rated as significantly more controllable than attributes generated in the Discrimination condition, $b = .91$, $SE = .03$, $t(1,848) = 34.29$, $p < .001$. Separating our analyses by question type, we found similar results: essay, $b = 1.04$, $SE = .04$, $t(874) = 28.50$, $p < .001$ and list, $b = .77$, $SE = .04$, $t(959) = 20.00$, $p < .001$.

We next specified a mixed-model regression, assessing the effects of condition on Phase 2 perceived relevance, including a random-intercept of item-generating participant. As expected, attributes generated in the Hiring condition were rated as significantly more relevant than attributes generated in the Discrimination condition, $b = 1.09$, $SE = .02$, $t(1,826) = 45.82$, $p < .001$. Again separating our analyses by question type, we found similar results: essay, $b = 1.19$, $SE = .03$, $t(863) = 35.82$, $p < .001$ and list, $b = 1.00$, $SE = .03$, $t(942) = 28.60$, $p < .001$.

Discussion

Using a preregistered converse logic approach, we find that when individuals were asked to think about discrimination, they listed attributes that were perceived as less controllable and relevant—as rated by a new set of participants—compared to when they were asked to think about hiring in general. Specifically, when asked to think about the attributes that might affect a presumably fair selection decision, people generated highly controllable and highly relevant attributes: bona-fide work qualifications like specialized skills and abilities. However, when prompted to think about an unfair, discriminatory decision, participants largely generated uncontrollable and irrelevant attributes, such as gender, sex, and race.

This pattern suggests that people may experience some myopia when thinking about selection decisions: Before being prompted to consider discrimination, they largely did not consider that

⁴ We standardized all Phase 1 responses by converting to lowercase letters, adding commas to separate lists of generated attributes, and correcting misspelled words. We then screened for duplicate responses (e.g., “race”), which gave us our final set of unique (nonduplicate) attributes to use as stimuli in Phase 2.

Table 3
Study 1a—Thematic Coding of Participants’ Open-Ended Responses

Theme	Description	Percent mentioned	Examples
Controllability	The extent to which people can have a direct say over the attribute they are being assessed on	13%	<p>“Sexual orientation is not a choice.”</p> <p>“People do not get to choose or affect their gender so telling someone they are unqualified because of something outside their control is unfair.”</p> <p>“If merit-based hiring is broadly speaking a benchmark for fairness, selection based on personal network connections in by definition unfair, as it disadvantages people for circumstances that are unrelated to their ability to succeed at the job, and that in many cases they may have no control over.”</p> <p>“Because a person really has no control over it.”</p>
Relevance	The extent to which the hiring attribute is relevant to job performance	54%	<p>“A person’s national origin doesn’t have anything to do with whether they can do a job well or not.”</p> <p>“Selecting someone based on socio-economic status is not relevant. It has no bearing on being qualified for a job.”</p> <p>“Because it shouldn’t have anything to do with whether a person is capable of doing their job or not.”</p> <p>“Because it has nothing to do with one’s merit or ability to do the job.”</p>

demographic characteristics might influence decisions to select a particular candidate over an equally competent other. Age represented the sole exception, which fits with recent work suggesting that people might perceive age as proxy for work experience (North, 2019). Also interestingly, many demographic attributes that are both objectively and legally the source of ongoing discrimination, such as caregiving responsibilities, were nevertheless rarely or *not* mentioned by participants as potential sources of discrimination. Importantly, these rarely mentioned attributes are those that might be construed as controllable and/or relevant. Notably, we find similar results in a pilot study that preceded Study 1b (see Online Supplemental Material). Overall, Study 1b suggests that perceptions of controllability and of relevance may relate to whether people find demographic attribute fair to use.

Study 2a

Study 2a aimed to quantitatively test whether perceived controllability and perceived relevance relate to people’s perceptions of fairness of selection based on a variety of demographic attributes, as we expected. Further, we expected to see stronger effects of perceived relevance than controllability on fairness. To examine these hypotheses, we recruited a large sample of U.S. adults, using a within-subjects design.

Method

Participants and Procedure

We recruited 1,000 new U.S. adults from Mechanical Turk to complete the study in exchange for \$1.00. We removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), resulting in a total of 813 participants, who made 15 ratings each ($N = 12,195$).

Participants read about and rated 15 selection attributes in a random order, following Study 1a: age, alma mater, caregiving responsibilities, disability, educational attainment, family origin (i.e., last name), gender, national origin, personal network connections, physical attractiveness, political affiliation, race, religious beliefs, sexual orientation, and socioeconomic status. After reading

the target attribute prompt, participants rated the perceived controllability, perceived relevance, and perceived fairness of using the attribute in selection decisions. Participants also completed two exploratory items, typicality and privacy, which did not load onto our theorized factors and thus were excluded (see Online Supplemental Material). Then, participants proceeded to the next target attribute. Finally, participants completed demographic information.

Measures

Perceived controllability ($r = .70$) and *perceived relevance* ($r = .77$) were measured as in Study 1b.

Perceived Fairness. To move beyond the single item used in Study 1a, we adapted a perceived fairness scale used by Hansen (1992; see also Reidenbach & Robin, 1990), using six items ($\alpha = .90$): “fair,” “justifiable,” “acceptable,” “immoral” (r), “harmful” (r), and “discriminatory” (r), rated on a scale from 1 (*not at all*) to 7 (*extremely*).

Results

As expected, factor analysis with principal axis factoring and varimax rotation supported a three-factor solution, revealing that each of our measured variables (perceived controllability, perceived relevance, and perceived fairness) assessed three unique constructs, which together explained 65% of the variance (see Online Supplemental Material). Additionally, see Table 5 for a summary of results across studies.

Effects of Controllability and Relevance

First, we specified a mixed-model regression assessing the effect of perceived controllability (centered) on perceived fairness, including a random-intercept effect of demographic attribute and participant. As expected, we found a significant effect of perceived controllability, $b = .22$, $SE = .01$, $t(10,890) = 25.84$, $p < .001$.

Second, in a similar model, we assessed the effect of perceived relevance (centered) on perceived fairness. As expected, we found a

Table 4
Correlations Among Variables (Studies 1b-4)

Variable	1	2	3	4	5	6	7	8	9	10	11	
1. Perceived controllability	—											
2. Perceived relevance	S1b: .50*** S2a: .51*** S2b: .44*** S2c: .46*** S3a: .21*** S3b: .10** S3c: .12** S4: .63***	—										
3. Perceived fairness	S2a: .69*** S2b: .63*** S2c: .62*** S3a: .38*** S3b: .24*** S3c: .30*** S4: .77***	S2a: .69*** S2b: .63*** S2c: .62*** S3a: .38*** S3b: .24*** S3c: .30*** S4: .77***	—									
4. Descriptive legality	S2b: .28*** S2b: .33*** S3a: .25*** S3a: .07	S2b: .42*** S2b: .59*** S3a: .66*** S3a: .18***	S2b: .54*** S2b: .74*** S3a: .60*** S3a: .31***	— S2b: .64*** — —	— — — —							
5. Prescriptive legality	S3a: .25*** S3a: .07	S3a: .66*** S3a: .18***	S3a: .60*** S3a: .31***	— —	— —	— —	— —	— —	— —	— —	— —	
6. Attribute curiosity	S3a: .07	S3a: .18***	S3a: .31***	—	—	—	—	—	—	—	—	
7. Attribute information seeking	S3a: .07	S3a: .18***	S3a: .31***	—	—	—	—	—	—	—	—	
8. Applicant rating pre-reveal	S3a: -.03	S3a: .06	S3a: .05	—	—	—	—	—	—	—	—	
9. Applicant rating post-reveal	S3a: -.08	S3a: -.19*** S4: .67***	S3a: -.19*** S4: .67***	— —	— —	— —	— —	— —	— —	— —	— —	
10. Perceived use	S4: .56***	S4: .67***	S4: .67***	—	—	—	—	—	—	—	—	
11. Psychological safety	S4: -.02	S4: -.03	S4: .00	—	—	—	—	—	—	—	—	
12. Job satisfaction	S4: -.01	S4: -.04	S4: -.02	—	—	—	—	—	—	—	—	

* $p < .05$. ** $p < .01$. *** $p < .001$.

Table 5*Effects of Perceived Controllability and Perceived Relevance on Perceived Fairness (Studies 2a, 2b, 2c, and 4)*

Variable (fixed)	Study 2a		Study 2b				Study 2c				Study 4					
	<i>b</i>	<i>SE</i>														
Intercept	2.68***	.17	2.68***	.10	2.66***	.18	2.66***	.09	2.68***	.17	2.68***	.10	3.14***	.48	3.14***	.37
Perceived controllability	.22***	.01			.32***	.03			.43***	.03			.27***	.02		
Perceived relevance			.52***	.01			.58***	.03			.57***	.03			.47***	.02
Variable (random)	<i>SD</i> ²															
Participant	.64		.41		—		—		—		—		.47		.44	
Demographic attribute	.44		.15		.47		.12		.41		.14		2.10		1.22	

*** $p < .001$.

significant effect of perceived relevance, $b = .52$, $SE = .01$, $t(10,500) = 64.72$, $p < .001$.

Next, we considered whether perceived relevance exerted a stronger effect than perceived controllability on fairness perceptions, using relative weights analysis (*rwa* package in R with 10,000 bootstraps; see also Tonidandel & LeBreton, 2011). The results yielded that both perceived controllability ($\epsilon_{raw} = .11$, 95% CI [.10, .11]) and perceived relevance ($\epsilon_{raw} = .39$, 95% CI [.37, .40]) accounted for a significant amount of the variance in perceived fairness. However, as expected, perceived relevance explained significantly more of the total variance than perceived controllability ($\Delta_{raw} = .28$, 95% CI [.26, .30], $p < .05$).

Discussion

Building on the findings of Studies 1a–b, Study 2a suggests that perceived controllability and perceived relevance relate to people's perceptions of the fairness of using demographic attributes in selection. Consistent with our theorizing, Study 2a additionally suggests that perceived relevance is a stronger predictor than perceived controllability. That is, rejecting an applicant due to a demographic attribute such as age may be perceived as fair, if age is perceived as relevant to the job, despite its uncontrollability by the individual, and despite its legally protected status.

Study 2b

Study 2b aimed to replicate Study 2a, with two changes. First, Study 2b employed a between-subjects design, to avoid potential participant fatigue. Second, we also measured perceived legality, to see if participants' impressions of fairness would persist above and beyond their understanding of legality. Study 2b was preregistered: <https://aspredicted.org/898mu.pdf>.

Method

Participants and Procedure

We recruited 1,500 new U.S. adults from Prolific to complete the study in exchange for \$.45. Following our preregistration, we removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), resulting in a final sample of 1,165 participants.

Study 2b followed the same procedure as Study 2a. However, in this study, participants rated only one demographic attribute and

additionally answered questions about the descriptive and prescriptive legality of using the given attribute in selection decisions.

Measures

Perceived relevance ($r = .66$), *perceived controllability* ($r = .60$), and *perceived fairness* ($\alpha = .92$) were measured as in Study 1b.

Descriptive Legality. We measured the extent to which participants believed selection based on the attribute is legal with the following item: "To what extent do you think that selection based on [attribute] is legal or illegal?" (1 = *definitely illegal*; 7 = *definitely legal*). Following our preregistration, we used this item as a control variable to see whether perceived controllability and relevance had an effect above and beyond knowledge of legality.

Prescriptive Legality. We also measured the extent to which participants think that selection based on the attribute *should* be legal with the following item: "To what extent do you think selection based on [attribute] should be legal or illegal?" (1 = *definitely should be illegal*; 7 = *definitely should be legal*). Following our preregistration, we used this item as an exploratory dependent variable.

Importantly, we again found that our measures supported our theorized factor structure, such that perceived relevance, perceived controllability, perceived fairness, and legality (both descriptive and prescriptive) each loaded onto unique factors (see Online Supplemental Material).

Results

Replicating Study 2a

First, we specified a mixed-model regression assessing the effect of perceived controllability (centered) on perceived fairness, including a random-intercept effect of demographic attribute. As expected, we found a significant effect of perceived controllability, $b = .32$, $SE = .03$, $t(1,152) = 10.69$, $p < .001$. Controlling for descriptive legality, the effect of perceived controllability remained significant, $b = .24$, $SE = .03$, $t(1,119) = 8.67$, $p < .001$.

Second, in a similar model, we assessed the effect of perceived relevance (centered) on perceived fairness. As expected, we found a significant effect of perceived relevance, $b = .58$, $SE = .03$, $t(1,083) = 21.72$, $p < .001$. Controlling for descriptive legality, the effect of perceived relevance remained significant, $b = .48$, $SE = .03$, $t(1,073) = 18.73$, $p < .001$.

Third, using relative weights analysis, we tested whether perceived relevance exerted a stronger effect than perceived

controllability on fairness perceptions. We confirmed that both perceived controllability ($\epsilon_{\text{raw}} = .05$, 95% CI [.03, .07]) and perceived relevance ($\epsilon_{\text{raw}} = .34$, 95% CI [.30, .39]) accounted for a significant amount of the variance in perceived fairness, but perceived relevance explained significantly more of the variance than perceived controllability ($\Delta_{\text{raw}} = .29$, 95% CI [.24, .34], $p < .05$), as expected.

Exploratory Analyses: Prescriptive Legality

Following our preregistration, we also explored how perceived controllability and perceived relevance relate to prescriptive legality. That is, to the degree individuals believe an attribute to be fair to use, then they may also believe that the attribute *should* be legal to use. Using a mixed-model regression, we found a significant effect of both perceived controllability (centered), $b = .10$, $SE = .03$, $t(924) = 3.79$, $p < .001$, and perceived relevance (centered), $b = .43$, $SE = .03$, $t(1,065) = 16.28$, $p < .001$, on prescriptive legality, in line with our theorizing. These effects remained significant, perceived controllability: $b = .05$, $SE = .02$, $t(357) = 2.54$, $p = .01$; perceived relevance: $b = .32$, $SE = .02$, $t(740) = 14.32$, $p < .001$, even after controlling for descriptive legality, $b = .46$, $SE = .02$, $t(803) = 20.38$, $p < .001$.

Exploratory Analyses: Attribute Clusters

Following our preregistration, we also considered whether some demographic attributes are considered more similar in terms of their perceived controllability and relevance than others. We conducted a two-dimensional cluster analysis with k -medoids, crossing perceived relevance (y-axis) and perceived controllability (x-axis) of attributes and using euclidean distance to calculate dissimilarity (*pam* function in R package “cluster”). Results revealed four clusters: (a) high

relevance/high controllability (e.g., educational attainment, caregiving); (b) low relevance/low controllability (e.g., gender, race); (c) ambiguous (e.g., religious beliefs, alma mater); and (d) high relevance/low controllability (e.g., age, disability; see Figure 1).

Finally, we explored the impact of this cluster structure on perceived fairness. We recoded the 15 attributes into four categories, representing the cluster structure. A one-way analysis of variance (ANOVA) revealed a significant difference by cluster type for perceived fairness, $F(3, 1,161) = 59.72$, $p < .001$, $\eta_p^2 = .13$. Specifically, each cluster was perceived as significantly different from the others regarding fairness ($p < .003$), with the high relevance/high controllability cluster perceived as most fair, followed by the high relevance/low controllability cluster, then the ambiguous cluster, and finally the low relevance/low controllability cluster. This pattern further suggests that perceived relevance may be a stronger predictor of fairness than is controllability.

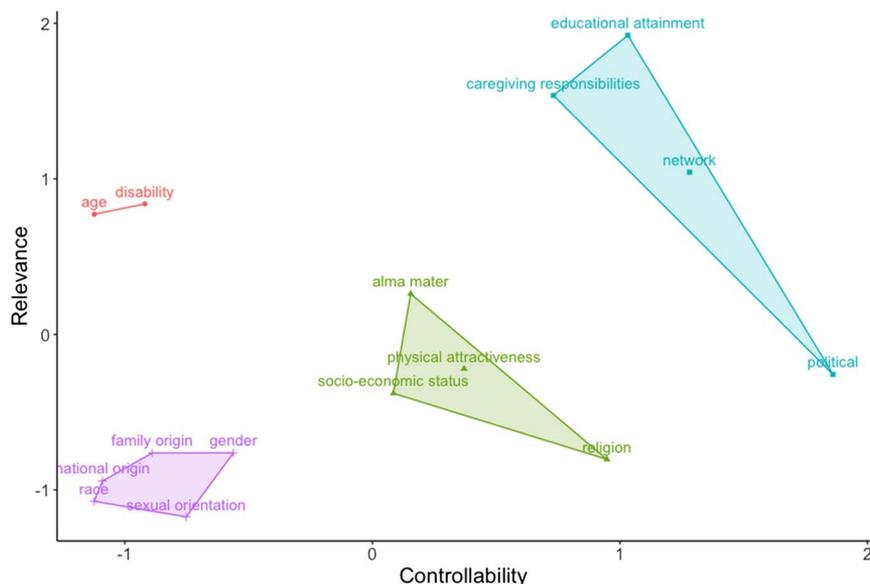
Exploratory Analyses: Group Membership

Given advantaged groups may be more likely to find discrimination to be fair (e.g., Major, Gramzow, et al., 2002; Purdie-Vaughns et al., 2008; see also Jetten, Iyer, et al., 2013), we considered whether the effects of perceived controllability and perceived relevance persisted above and beyond participants' own group membership.

First, we specified a mixed-model regression assessing the effect of perceived controllability, perceived relevance, and age (each centered) on perceived fairness, including a random-intercept effect of demographic attribute. Results revealed a significant effect of perceived controllability, $b = .08$, $SE = .03$, $t(944) = 2.61$, $p = .009$, and perceived relevance, $b = .54$, $SE = .03$, $t(1,075) = 18.17$, $p < .001$, but no effect of age ($p = .35$).

Next, we tested race (1 = White; -1 = Non-White) in a similar model. We found a significant effect of perceived controllability,

Figure 1
Study 2b—Two-Dimensional Cluster Analysis



Note. See the online article for the color version of this figure.

$b = .08$, $SE = .03$, $t(940) = 2.67$, $p = .008$, and perceived relevance, $b = .54$, $SE = .03$, $t(1,073) = 18.19$, $p < .001$, but no effect of race ($p = .35$). We also considered participant race coded as a six-category factor (reference = White; 1 = Asian, 2 = Black, 3 = Latino, 4 = Native, 5 = Other). We again found a significant effect of perceived controllability, $b = .08$, $SE = .03$, $t(933) = 2.62$, $p = .009$, and perceived relevance, $b = .54$, $SE = .03$, $t(1,068) = 18.15$, $p < .001$, but no effect of race ($p > .24$).

Finally, we ran a similar model testing the effect of gender (1 = Male; -1 = Female or Nonbinary). We again found a significant effect of perceived controllability, $b = .07$, $SE = .03$, $t(952) = 2.54$, $p = .01$, and perceived relevance, $b = .53$, $SE = .03$, $t(1,077) = 17.94$, $p < .001$. This time, we also found a significant effect of gender, $b = .11$, $SE = .03$, $t(1,151) = 3.30$, $p = .001$, such that men were more likely to perceive selection across demographic attributes to be fair. Nevertheless, the effects of perceived controllability and relevance persisted.

We also explored the same mixed-model regression with all controls simultaneously. As expected, we found a significant effect of perceived controllability, $b = .07$, $SE = .03$, $t(949) = 2.50$, $p = .01$, and perceived relevance, $b = .53$, $SE = .03$, $t(1,074) = 17.91$, $p < .001$. We also found a significant effect of gender, $b = .11$, $SE = .03$, $t(1,148) = 3.36$, $p < .001$, but no effect of age ($p = .33$) nor race (binary coding; $p = .53$). Notably, we also find similar results across studies (see Online Supplemental Material).

Discussion

Study 2b offered a preregistered replication of the results from Study 2a: Perceived controllability and perceived relevance both relate to people's perceptions of fairness of demographic-based selection, across a wide range of such attributes. Again, we find that perceived relevance exerts more influence.

We also find that the effects of perceived controllability and perceived relevance on fairness persist above and beyond perceived legality. Relatedly, individuals' own perceptions of controllability and relevance affect their beliefs about whether a given attribute *should* be legal or illegal, above and beyond their perceptions of current legality. Interestingly, we also find that perceptions of controllability and relevance remain significant predictors above and beyond participants' own demographic group membership (see also General Discussion section). Overall, these results support our suggestion that discrimination may occur because some attributes (albeit illegal) are perceived as fair to use, as determined by controllability and relevance principles of meritocracy.

Further, Study 2b provides evidence that demographic attributes may cluster into four perceived categories: clearly unfair (low relevance, low controllability), clearly fair (high relevance, high controllability), and two somewhat fair (ambiguous cluster; high relevance, low controllability cluster). We find that several legally protected attributes, including age, disability, and caregiving responsibilities, fall into the somewhat fair and even fair clusters. Notably, we also replicate this overall cluster pattern in Study 2a (see Online Supplemental Material). The cluster pattern again suggests that individuals may value perceived relevance in particular, above and beyond the law.

Study 2c

Next, we aimed to replicate our results specifically among a sample of relative experts: individuals who make hiring decisions

regularly as part of their jobs. Study 2c used the same procedure as Study 2b.

Method

Participants

We recruited 1,500 U.S. adults from Prolific, who had previously answered affirmatively to the prescreening question "Do you have any experience in making hiring decisions (i.e., have you been responsible for hiring job candidates)?," and who then completed the study in exchange for \$.45. We removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), resulting in a final sample of 1,056 participants.

Measures

Perceived relevance ($r = .67$), *perceived controllability* ($r = .49$), and *perceived fairness* ($\alpha = .93$) were measured as in Study 1b.

Results

First, we specified a mixed-model regression assessing the effect of perceived controllability (centered) on perceived fairness, including a random-intercept effect of demographic attribute. As expected, we found a significant effect of perceived controllability, $b = .43$, $SE = .03$, $t(1,038) = 13.42$, $p < .001$.

Second, in a similar model, we assessed the effect of perceived relevance (centered) on perceived fairness. As expected, we found a significant effect of perceived relevance, $b = .57$, $SE = .03$, $t(1,011) = 20.81$, $p < .001$.

Third, using relative weights analysis, we tested whether perceived relevance exerted a stronger effect than perceived controllability on fairness perceptions. We confirmed that both perceived controllability ($e_{\text{raw}} = .10$, 95% CI [.07, .13]) and perceived relevance ($e_{\text{raw}} = .31$, 95% CI [.25, .36]) accounted for a significant amount of the variance in perceived fairness, but perceived relevance explained significantly more of the variance than perceived controllability ($\Delta_{\text{raw}} = .21$, 95% CI [.14, .27], $p < .05$), as expected.

Discussion

Study 2c again replicated our earlier results, demonstrating similar effects even among individuals who make hiring decisions regularly. Perceived controllability and perceived relevance both relate to perceptions of fairness of demographic-based selection, again across a wide range of attributes. We again find that perceived relevance exerts more influence, even among relative experts in the hiring process.

Notably, even those who do not make formal hiring decisions regularly as part of their current jobs may still observe discrimination and make reporting decisions, and they may make discriminatory decisions directly. For instance, individuals make daily decisions about such varied things as where to shop, to whom to speak, to whom to award the project, or whom to hire as a babysitter—each of which has the potential to be discriminatory. Altogether, Studies 2a–c offer evidence that selection decisions based on given attributes may result in discrimination not in spite of

the perceived unfairness of the attributes, but *because* these attributes are perceived as fair.

Study 3a

In Study 3a, we aimed to establish causality by experimentally manipulating controllability and relevance during a hiring simulation. We theorized that to the extent that these two principles influence perceived fairness of a demographic attribute, then they should also affect people's willingness to seek and use information about the given attribute during the selection process. Therefore, Study 3a includes behavioral measures of people's willingness to seek and ultimately use information about demographic attributes when making a selection decision. We used a preregistered, 2 (high vs. low relevance) \times 2 (high vs. low controllability) between-subjects design: <https://aspredicted.org/3q2u.pdf>.

Method

Participants

We recruited 500 new U.S. adults from Prolific to complete the study in exchange for \$.75. Following our preregistration, we removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), resulting in a final sample of 459 participants.

Procedure

To ensure people's preexisting beliefs about an attribute did not interfere with the manipulations (e.g., Hegarty & Golden, 2008), we relied on a demographic attribute about which U.S. participants did not have preexisting beliefs—"sortable." Participants read:

Sortable is a French word used to describe people sometimes. This is a word that has no direct English translation. In France, people are described as sortable when they are very understated and proper, and they tend to avoid too much attention or embarrassing moments in public. Sortable helps describe this tendency people may or may not have.

Participants then read about the controllability and relevance (counterbalanced) of the attribute and answered manipulation check questions. Next, they reported their perceived fairness of using sortable in selection decisions, then indicated the extent to which they would be curious to learn about job candidates' sortable when making selection decisions.

Then, participants entered a hiring simulation, in which they learned they would be reviewing a job candidate's application for a consulting position at a company. After reviewing the applicant's cover letter, participants were asked to evaluate the applicant. Then, participants were offered the chance to choose what type of information they wanted to further gather about the candidate from their full application. Here, we tested whether participants opted to see more information about the candidate's *sortable* (i.e., the target attribute). Regardless of whether they selected to learn more information about the candidate's *sortable*, we showed participants all available application information, including stereotypically negative information about the target attribute: "not good—not sortable at all" (see full materials in Online Supplemental Material). Finally, participants were asked to evaluate the applicant a second time, after

learning this negative information. Participants finished the experiment by reporting their demographics.⁵

Measures

Independent Variables.

Controllability. Participants were randomly assigned to either the High ($n = 224$) or the Low Controllability ($n = 235$) condition. Participants read that "In general, French people use the word *sortable* to describe a tendency in people that [is/is not] under their own choice or control. In other words, a person [can/has no] control [over] their *sortable*. [People are either born *sortable*, or they are not./Even if someone started out one way, they can choose to shift their *sortable* over time.]"⁶

Relevance. Participants were randomly assigned to either the High ($n = 231$) or the Low Relevance ($n = 228$) condition. Participants read that

In general, most job recruiters in French companies consider *sortable* [relevant/irrelevant]. That is, companies know that how *sortable* a person seems [ultimately/does not] affect[s] their job performance [a lot/whatever]. Therefore, recruiters often [pay attention to/ignore] information about job candidates' *sortable* when making a hiring decision.

Dependent Variables. *Perceived controllability* ($r = .82$), *perceived relevance* ($r = .85$), and *perceived fairness* ($\alpha = .88$) were measured as in Study 2a.

Attribute Curiosity. We measured participants' curiosity regarding the manipulated attribute with the following four items: "I am curious to know what the applicant's *sortable* is," "The applicant's *sortable* should influence my hiring decision," "It is justifiable to ask about a person's *sortable* on a job application," and "The applicant's *sortable* will influence my hiring decision" (1 = *not at all*, 7 = *extremely*; $\alpha = .92$).

Attribute Information Seeking. We measured whether participants would seek information about the target demographic attribute by having them select whether or not they wanted to learn more about this attribute. Participants could learn more information about the job candidate's *sortable* by clicking on "Interviewer's notes on applicant: Other Observations (comment on *sortable*)." Specifically, information about the target attribute was buried within a list of 10 additional distractor items (e.g., additional skills, hobbies, awards and honors, interviewer's notes on applicant's motivation, notes on strengths, notes on weaknesses). We asked: "Please indicate which

⁵ In demographics, we asked whether participants were fluent French speakers and found that only five reported fluency in Study 3a and five again in Studies 3b and 3c. Following our preregistration, we replicated our analyses excluding fluent participants and found that the pattern and significance of results persisted in both Studies 3a–c.

⁶ In a separate pretest of these manipulations, we found that both were effective (see Online Supplemental Material). In the current experiment, we also found that the manipulations were successful, as expected: Individuals in the High Controllability condition perceived *sortable* as significantly more controllable ($M = 5.06$, $SD = 1.23$) than those in the Low Controllability condition, $M = 1.80$, $SD = 1.35$, $t(457) = 26.97$, $p < .001$, $d = 2.52$. Individuals in the High Relevance condition perceived *sortable* as significantly more relevant ($M = 5.08$, $SD = 1.30$) than those in the Low Relevance condition, $M = 1.89$, $SD = 1.41$, $t(457) = 25.20$, $p < .001$, $d = 2.35$. We found the same results in Study 3b and Study 3c (see Online Supplemental Material). Finally, and following our preregistration, we replicated our main text analyses controlling for manipulation counterbalance order in Studies 3a–c. Both the pattern and significance of our results persisted.

information about the applicant you would like to see in order to help you evaluate the applicant." Participants were allowed to select as many boxes containing information about the applicant as they wanted. Selecting the target box was coded as 1, and not selecting this information was coded as 0.

Applicant Ratings. Before and after we revealed negative information about the candidate's *sortable*, we measured participants' evaluations of the applicant with the following five items (adapted from Uhlmann & Cohen, 2005): "The applicant is highly motivated," "The applicant should be invited for a final interview," "The applicant is a strong candidate," "The applicant will be successful at the job," and "The applicant should be hired" ($\alpha = .93$ pre; $\alpha = .93$ post). These items were rated on a scale from 1 (*not at all*) to 7 (*extremely*).

Results

In the following analyses, both the Controllability and Relevance conditions were contrast coded (1 = High, -1 = Low).

Perceived Fairness

Using an independent samples *t*-test, those in the High Controllability condition perceived *sortable* as more fair to use in selection ($M = 4.49$, $SD = 1.30$) than those in the Low Controllability condition, $M = 3.75$, $SD = 1.44$, $t(456) = 5.81$, $p < .001$, $d = .54$.

In another *t*-test, individuals in the High Relevance condition perceived *sortable* as more fair to use in the selection process ($M = 4.51$, $SD = 1.39$) than those in the Low Relevance condition, $M = 3.70$, $SD = 1.35$, $t(457) = 6.31$, $p < .001$, $d = .59$.

Further, relative weights analysis yielded that both Controllability condition ($\epsilon_{\text{raw}} = .06$, 95% CI [.02, .10]) and Relevance condition ($\epsilon_{\text{raw}} = .08$, 95% CI [.03, .12]) accounted for a significant amount of the variance in perceived fairness. Although Relevance condition explained more variance than Controllability condition descriptively, they were statistically similar ($\Delta_{\text{raw}} = .02$, 95% CI [-.05, .08]).

Attribute Curiosity

We expected that perceptions of controllability and relevance should increase participants' perceptions of fairness of using the attribute, which in turn should increase their curiosity about that attribute in selection decisions.

Using a *t*-test, we found that those in the High Controllability condition were more curious to learn about *sortable* in the selection process ($M = 3.18$, $SD = 1.60$) than those in the Low Controllability condition, $M = 2.72$, $SD = 1.57$, $t(455) = 3.09$, $p = .002$, $d = .29$. Similarly, individuals in the High Relevance condition were more curious to learn about *sortable* in the selection process ($M = 3.78$, $SD = 1.47$) than those in the Low Relevance condition, $M = 2.11$, $SD = 1.24$, $t(446) = 13.14$, $p < .001$, $d = 1.23$.

We also tested whether there was an indirect effect of Controllability condition on attribute curiosity, through perceived fairness. The mediation model (5,000 bootstraps) revealed a significant indirect effect ($b = .25$, $SE = .05$, 95% CI [.16, .35]; Figure 2a). We then tested whether there was an indirect effect of Relevance condition on attribute curiosity, through perceived fairness. We again found a significant indirect effect ($b = .22$, $SE = .04$, 95% CI [.15, .31]; Figure 2b).

Attribute Information Seeking

We also expected that perceptions of controllability and relevance should increase participants' willingness to seek out information about that attribute in selection decisions.

To test this hypothesis, we specified a binomial regression, assessing the effect of Controllability condition and Relevance condition on attribute information seeking. We did not find an effect of Controllability condition, $b = .08$, $SE = .10$, $z(458) = .88$, $p = .38$. However, we found a significant effect of Relevance condition, $b = .49$, $SE = .10$, $z(458) = 5.07$, $p < .001$, as expected.

Next, we tested whether there was an indirect effect of Controllability condition on attribute information seeking, through perceived fairness. Despite the lack of main effect of Controllability condition on information seeking, the mediation model (5,000 bootstraps) revealed a significant indirect effect as expected ($b = .04$, $SE = .01$, 95% CI [.02, .06]; Figure 2c). We then tested whether there was an indirect effect of Relevance condition on attribute information seeking, through perceived fairness. We again found a significant indirect effect ($b = .04$, $SE = .01$, 95% CI [.02, .06]; Figure 2d).

Applicant Ratings

We expected that perceptions of controllability and relevance should increase participants' use of the attribute in their evaluations. Specifically, we expected that negative information about that attribute should affect participants' ratings of the job applicant more when they found that attribute fair to use.

To test this hypothesis, we specified a regression assessing the effect of perceived fairness on post-reveal ratings of the applicant (i.e., after negative information was revealed). We found that perceived fairness had a significant effect on post-reveal ratings, as expected, $b = -.17$, $SE = .04$, $t(457) = -4.12$, $p < .001$. Using the same model to instead predict pre-reveal ratings of the applicant (*before* the negative information was revealed), we found no effect of perceived fairness, $b = .04$, $SE = .03$, $t(457) = 1.12$, $p = .27$, as expected. This suggests that when negative information about the target attribute was revealed, participants only *used* this information in their ratings to the extent they found the target attribute fair to use in selection.

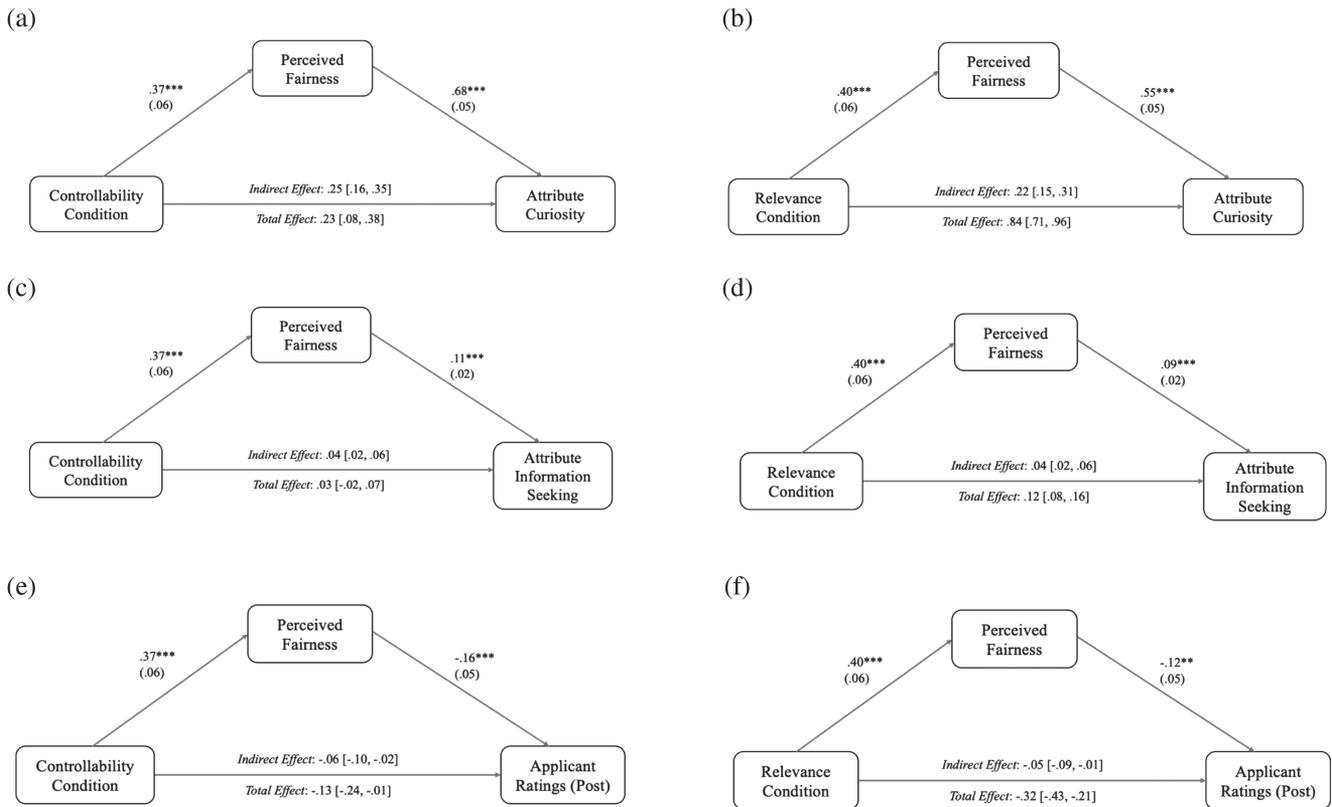
Next, we tested whether there was an indirect effect of Controllability condition on post-reveal ratings, through perceived fairness. The mediation model (5,000 bootstraps) revealed a significant indirect effect ($b = -.06$, $SE = .02$, 95% CI [-.10, -.02], Figure 2e). We then tested whether there was an indirect effect of Relevance condition on post-reveal ratings, through perceived fairness, and we again found a significant indirect effect ($b = -.05$, $SE = .02$, 95% CI [-.09, -.01]; Figure 2f).

Discussion

Using a preregistered experimental design, Study 3a directly manipulated perceptions of controllability and relevance of an attribute, which in turn affected perceived fairness of using that attribute in a hiring simulation. As such, Study 3a replicated the pattern of Studies 1a, 1b, 2a, and 2b and lends experimental support to our theorizing.

Further, we find that controllability and relevance in turn affect people's *use* of demographic attributes during the selection process, via

Figure 2
Study 3a—Mediation Analyses



** $p < .01$. *** $p < .001$.

fairness perceptions. Specifically, when we led people to believe an attribute was controllable and relevant, people were more curious to learn about the attribute, and they actually *sought* information regarding that attribute during the hiring process. Given the availability of notes regarding personal details of an applicant may be unexpected among our U.S.-based participants, and given there was no penalty for seeking the information, people might have been inherently quite curious—working against our hypotheses, and thus providing a conservative test. Nevertheless, as predicted, we find experimental evidence that controllability and relevance shape people’s curiosity and behavior regarding the attribute.

Finally, controllability and relevance ultimately affected ratings of applicants who were revealed to have undesirable levels of the target attribute: The more the attribute was perceived to be controllable and relevant, the more negatively the applicant was rated, via perceived fairness. As such, we find evidence that people’s reliance on meritocratic principles of controllability and relevance might ultimately lead them to use demographic attributes in selection and discriminate.

Study 3b

To replicate and extend Study 3a, we next aimed to explore whether individuals’ perceptions changed when exposed to *explicit* information about the illegality of demographic attributes. Study 3b

used the same 2 (high vs. low relevance) \times 2 (high vs. low controllability) between-subjects design.

Method

Participants

We recruited 400 new U.S. adults from Prolific to complete the study in exchange for \$.45. We removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), resulting in a final sample of 361 participants.

Procedure

Study 3b used the same hiring scenario as in Study 3a. However, in Study 3b before the manipulations, we explicitly specified that selecting based on *sortable* (the target demographic attribute) was illegal: “In French society, hiring someone based on their sortable has been banned. That is, it is illegal to hire someone based on sortable alone.”

Measures

Independent Variables.

Controllability. Participants were randomly assigned to either the High ($n = 174$) or the Low Controllability ($n = 187$) condition, as in Study 3a.

Relevance. Participants were randomly assigned to either the High ($n = 191$) or the Low Relevance ($n = 170$) condition, as in Study 3a.

Dependent Variables. *Perceived controllability* ($r = .86$), *perceived relevance* ($r = .86$), and *perceived fairness* ($\alpha = .88$) were measured as in Study 3a.

Comprehension Check. We confirmed that participants understood the illegality instruction with the single item: “Based on what you read earlier, is selection based on sortable illegal or legal in French society?” (1 = *illegal*; 2 = *not sure*; 3 = *legal*). Two hundred and ninety-eight (83%) participants responded to this question correctly (response = 1). We retain all participants here, but the pattern and significance of results persisted when we remove participants who failed the comprehension check (see Online Supplemental Material).

Results

In the following analyses, both the Controllability and Relevance conditions were contrast coded (1 = High, -1 = Low).

First, an independent samples t -test showed that individuals in the High Controllability condition perceived sortable as more fair to use in the selection process ($M = 3.86$, $SD = 1.30$) than those in the Low Controllability condition, $M = 3.29$, $SD = 1.39$, $t(359) = 4.05$, $p < .001$, $d = .43$.

In another t -test, we found that individuals in the High Relevance condition perceived sortable as more fair ($M = 3.76$, $SD = 1.39$) than those in the Low Relevance condition, $M = 3.35$, $SD = 1.33$, $t(357) = 2.81$, $p = .005$, $d = .30$.

Finally, relative weights analysis revealed that both Controllability condition ($\epsilon_{\text{raw}} = .04$, 95% CI [.001, .08]) and Relevance condition ($\epsilon_{\text{raw}} = .02$, 95% CI [-.01, .05]) accounted for a significant or marginally significant amount of the variance in perceived fairness. However, both conditions yielded statistically similar effects ($\Delta_{\text{raw}} = -.02$, 95% CI [-.07, .03]).

Discussion

Study 3b further replicated the results found in Study 3a: Perceptions of both controllability and relevance of a demographic attribute drive perceptions of the fairness of using that attribute in selection. Further, Study 3b demonstrates that even in the presence of explicit information about the illegality of an attribute, perceptions of controllability and relevance still affect fairness judgments. These findings provide further evidence that principles of meritocracy, above and beyond legality, can contribute to perceptions of fairness in selection, thus potentially perpetuating discrimination.

Study 3c

In Study 3c, we aimed to replicate the results from Study 3a, utilizing even more conservative manipulations—specifically removing behavioral norms information. Study 3c otherwise employed the same 2 (high vs. low relevance) \times 2 (high vs. low controllability) between-subjects experimental design as Study 3a.

Method

Participants

We recruited 500 new U.S. adults from Prolific to complete the study in exchange for \$.45. We removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), resulting in a final sample of 424 participants.

Measures

Independent Variables.

Controllability. Participants were randomly assigned to either the High ($n = 217$) or the Low Controllability ($n = 207$) condition, as in Study 3a.

Relevance. Participants were randomly assigned to either the High ($n = 206$) or the Low Relevance ($n = 218$) condition. Participants read a manipulation focusing only on the factual basis for relevance, without recruiters' normative behavior: “In general, job recruiters in French companies consider sortable [relevant/irrelevant]. That is, companies know that how sortable a person is [ultimately/does not] affect[s] their job performance [a lot/whatsoever].”

Dependent Variables. *Perceived controllability* ($r = .78$), *perceived relevance* ($r = .75$), and *perceived fairness* ($\alpha = .87$) were measured as in Study 2a.

Results and Discussion

Using an independent samples t -test, those in the High Controllability condition perceived sortable as more fair to use in selection ($M = 4.43$, $SD = 1.34$) than those in the Low Controllability condition, $M = 4.04$, $SD = 1.42$, $t(417) = 2.86$, $p = .004$, $d = .28$, as expected.

In another t -test, individuals in the High Relevance condition perceived sortable as more fair to use in the selection process ($M = 4.54$, $SD = 1.22$) than those in the Low Relevance condition, $M = 3.95$, $SD = 1.48$, $t(414) = 4.52$, $p < .001$, $d = .44$, as expected.

Further, relative weights analysis yielded that both Controllability condition ($\epsilon_{\text{raw}} = .02$, 95% CI [-.01, .04]) and Relevance condition ($\epsilon_{\text{raw}} = .04$, 95% CI [.01, .08]) accounted for a significant or marginally significant amount of the variance in perceived fairness. Although Relevance condition explained more variance than Controllability condition descriptively, they were statistically similar ($\Delta_{\text{raw}} = .02$, 95% CI [-.02, .07]).

Overall, the results of Study 3c again replicated those of Study 3a: The controllability and relevance of a given attribute affected perceived fairness of using that attribute.⁷

⁷ In addition to Studies 3a–c, we conducted another experiment aiming to manipulate participants' perceptions of controllability and relevance for a variety of familiar demographic attributes, about which participants may have preexisting beliefs (e.g., age, disability, religion; Crandall & Eshleman, 2003; Hegarty & Golden, 2008). Even in this context of preexisting beliefs, we replicate the indirect effect patterns from Studies 3a–c: manipulated controllability and manipulated relevance affect perceptions, which in turn affect perceived fairness and discriminatory behavior (see Supplemental Study 3a in Online Supplemental Material).

Study 4

Finally, we sought to investigate potential downstream consequences of perceiving selection decisions to be based on (un)controllable and (ir)relevant demographic attributes in the real world. Indeed, previous work has suggested that perceived discrimination in organizations can depress both psychological safety and satisfaction (e.g., Sanchez & Brock, 1996; Schmitt et al., 2014). In contrast, people are more satisfied with and willing to support organizations they perceive as using fair criteria (e.g., Brockner et al., 1994). Drawing on these insights, we tested whether controllability and relevance may act as antecedents to such outcomes. Specifically, we aimed to test whether employees' perceptions of their own organization's use of selection attributes that are controllable and/or relevant in turn are associated with employees' psychological safety and job satisfaction, through perceived fairness. Study 4 used a correlational, stimulus-sampling design to survey working adults about their organizational experience.

Method

Participants

We recruited 500 new, employed U.S. adults from Mechanical Turk to complete a survey in exchange for \$1.00. We removed incomplete and duplicate responses (based on IP address, location coordinates, and platform-based unique participant identifier), as well as those not located in the U.S. (based on IP address), resulting in a final sample of 300 participants, who made nine ratings each ($N = 2,700$).

Procedure

First, participants each rated nine selection attributes: educational attainment, past work performance, years of work experience, gender, race, sexual orientation, age, disability, and personal network connections. In this study, we included bona-fide work qualifications along with demographic attributes to offer a variety of expected levels of fairness. Next, participants were asked to reflect on their current organization and rated their own psychological safety and satisfaction experience. Finally, participants reported demographics.

Measures

Independent Variables.

Perceived Relevance. This was measured with the item "To what extent are the attributes below relevant to most jobs," asking participants to rate each of the nine attributes on a scale from 1 (*not at all relevant*) to 7 (*extremely relevant*).

Perceived Controllability. This was measured with the item "To what extent are the attributes below within someone's personal control" asking participants to rate each of the nine attributes on a scale from 1 (*not at all within personal control*) to 7 (*extremely within personal control*).

Perceived Use of Attributes. This was measured with the item "How much do these attributes below influence hiring or promotion in your current organization/workplace?" asking participants to rate each of the nine attributes on a scale from 1 (*no influence*) to 7 (*extreme influence*).

Dependent Variables.

Perceived Fairness. We measured participants' perceptions of fairness of each of the nine attributes with a single item: "To what extent are the attributes below a fair way of selecting employees for hiring or promotion?" asking participants to rate each of the nine attributes on a scale from 1 (*not at all fair*) to 7 (*extremely fair*).

Psychological Safety. Psychological safety was measured with nine items, including the 7-item scale used by Edmondson (1999; 1 = *not at all*, 7 = *extremely*) and two additional items: "I feel supported at my organization" and "I am able to freely express my opinions and views at my organization" ($\alpha = .88$).

Job Satisfaction. We measured job satisfaction using seven items, including four items adapted from Spector (1994; 1 = *not at all*, 7 = *extremely*) and three additional items: "I experience a high level of conflict at my job" (reversed), "My job causes me a lot of stress" (reversed), and "I am doing well at my job" ($\alpha = .82$).⁸

Results

Replicating Studies 2a–3c

We used a mixed-model approach, accounting for random-intercept effects of attribute and of participant (see Table 4). First, we assessed the effect of perceived controllability on perceived fairness. As expected, we found a significant effect of perceived controllability, $b = .27$, $SE = .02$, $t(2,696) = 15.40$, $p < .001$.

Next, we assessed the effect of perceived relevance on perceived fairness. As expected, we found a significant effect of perceived relevance, $b = .47$, $SE = .02$, $t(2,694) = 28.79$, $p < .001$.

Relative weights analysis again revealed that perceived relevance ($\epsilon_{\text{raw}} = .40$, 95% CI [.37, .42]) explained more of the variance in perceived fairness than did perceived controllability ($\epsilon_{\text{raw}} = .24$, 95% CI [.22, .26]), $\Delta_{\text{raw}} = .16$, 95% CI [.12, .20], $p < .05$.

Downstream Consequences

We expected to find that as perceived fairness decreases, the effect of perceived use on outcomes (psychological safety, job satisfaction) should become more negative. Further, larger discrepancies between perceived fairness of and perceived use of demographic attributes in selection should predict more negative outcomes. Given the structure of our data (assessing perceptions of multiple attributes within only one place of employment per participant), we employed mixed-effect modeling (*lmer* in R). We assessed fixed effects of perceived fairness of using the target demographic attribute in selection (centered), perceived use of the target demographic attribute in selection (centered), and their interaction on the outcome variables, including a random-intercept effect of attribute (see Table 6). Although our primary dependent variables were measured at the level of participant, this multilevel approach allows for an improved understanding of the relationship

⁸ This survey also included four additional downstream consequences: organizational fairness, organizational justice, organizational commitment, and organizational citizenship behaviors. Given both space constraints and the high correlations among these additional constructs (r 's range: .20–.77), we report these results and single item privacy in Online Supplemental Material. Overall, results suggested similar effects as reported in the main text; as perceived fairness diminishes (via perceived controllability and relevance), perceived use of attributes predicts worse employee outcomes.

Table 6

Study 4—Effects of Perceived Use and Perceived Fairness on Employee Outcomes (Interaction)

Variable (fixed)	Psychological safety		Job satisfaction	
	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>
Intercept	4.50***	.04	4.77***	.04
Perceived fairness	.04*	.02	.004	.01
Perceived use	-.15***	.01	-.12***	.01
Perceived fairness × Perceived use	.07***	.01	.06***	.01
Variable (random)	<i>SD</i> ²		<i>SD</i> ²	
Demographic attribute	.10		.01	

* $p < .05$. *** $p < .001$.

between outcomes (measured only once per participant rather than repeated at the attribute level) and perceived fairness and use (repeated at the attribute level), given the structure of our data (Croon & van Veldhoven, 2007).⁹

Psychological Safety. We found a significant main effect of perceived use, $b = -.15$, $SE = .01$, $t(2,436) = -10.38$, $p < .001$, and perceived fairness, $b = .04$, $SE = .02$, $t(183) = 2.35$, $p = .02$. As predicted, these effects were qualified by a significant interaction, $b = .07$, $SE = .01$, $t(2,369) = 12.36$, $p < .001$.

We decomposed the interaction using the Johnson–Neyman technique. When perceived fairness of attributes was less than -1.72 (centered), perceptions of organizational use of these attributes were *negatively* associated with participants' psychological safety at work (e.g., overutilized). However, when perceived fairness of attributes was greater than 2.79 (centered), perceptions of organizational use of these attributes were *positively* associated with participants' psychological safety at work. Decomposed differently, when perceived organizational use of attributes was less than -1.04 (centered), perceptions of attribute fairness were negatively associated with participants' psychological safety at work (e.g., underutilized). When perceived organizational use of attributes was greater than $-.09$ (centered), perceived attribute fairness was positively associated with participants' psychological safety.

Job Satisfaction. We specified a similar mixed-effect model as above and found a significant main effect of perceived use, $b = -.12$, $SE = .01$, $t(2,424) = -8.37$, $p < .001$, but no effect of perceived fairness, $b = .004$, $SE = .01$, $t(168) = .30$, $p = .76$, when predicting job satisfaction. However, as predicted, we found a significant interaction, $b = .06$, $SE = .01$, $t(2,369) = 11.08$, $p < .001$.

Decomposing the interaction, we found that when perceived fairness of attributes was less than 1.47 (centered), perceptions of organizational use of these attributes were negatively associated with participants' job satisfaction (e.g., overutilized). However, as expected, when perceived fairness of attributes was greater than 2.62 (centered), perceptions of organizational use of these attributes were positively associated with participants' job satisfaction. Decomposed differently, when perceived organizational use of attributes was less than $-.61$ (centered), perceptions of attribute fairness were negatively associated with participants' job satisfaction (e.g., underutilized). When perceived organizational use of attributes was greater than $.41$ (centered), perceptions of attribute fairness were positively associated with participants' job satisfaction.

Discussion

Replicating our previous studies, Study 4 again demonstrates that perceived controllability and relevance of demographic attributes relate to perceived fairness of using those attributes in selection. Moreover, Study 4 also demonstrates that these perceptions relate to important downstream outcomes among working adults: psychological safety and job satisfaction. We find that perceived fairness of attributes (undergirded by perceived controllability and relevance) interacts with perceived use of attributes such that employees felt safer and more satisfied with their organizations when perceived use and perceived fairness of attributes are more congruent, as compared to incongruent. In particular, employees' perceiving their organizations to overutilize attributes (low fairness—high use) may be more damaging than underutilizing attributes (high fairness—low use).

While we find a similar pattern across a range of attributes, we also again find evidence of clustering. For instance, highly uncontrollable and irrelevant attributes (e.g., gender, sexual orientation), and separately, highly controllable and relevant attributes (e.g., education, past performance) cluster together, yielding similar interactive effects. Future work might build on these results, and the cluster analyses in Studies 2a–2b, to explore how perceived relevance and controllability may lead some attributes to be considered similar to one another. For instance, just as different forms of stigma cluster along perceived shared dimensions (Pachankis et al., 2018), demographic attributes more generally may cluster to the extent they are perceived to share core dimensions important to meritocracy.

Overall, Study 4 offers correlational but real-world data, which suggests that individuals feel safe and satisfied when their organizations select employees on the basis of attributes that they *perceive* as controllable and relevant. As such, to the extent employees find some discriminatory attributes fair to use (e.g., age, disability), then they might remain satisfied with and support organizations that in fact engage in discrimination.

General Discussion

We proposed that people evaluate the fairness of demographic-based selection using two core principles of merit: controllability and relevance. Across nine studies, using both qualitative and quantitative approaches, we find evidence supporting this hypothesis. Further, perceptions of relevance more strongly affected perceived fairness than did perceptions of controllability. This suggests that some demographic attributes, such as age, disability, or caregiving responsibilities, may be perceived as fair to use in selection because people perceive them to be relevant to some jobs, even though they are uncontrollable and legally protected.

Perceived controllability and relevance, as drivers of fairness perceptions, also yield important downstream consequences. For instance, Study 3a demonstrates behavioral consequences: When people are led to perceive a demographic attribute as controllable and relevant, and therefore fair, they are more willing to seek and use information about the given attribute in selection decisions. Similarly, we find that perceived controllability and relevance shape

⁹ We supplemented this approach by also considering whether the difference score of perceived fairness and perceived use within each attribute predicted downstream consequences. The pattern and significance of results remained (see Online Supplemental Material).

perceptions of fairness of demographic attributes, which when coupled with perceptions of the actual use of those attributes, relate to employees' psychological safety and job satisfaction. As such, discrimination may persist, not in spite of people's beliefs that such discrimination is unfair, but because they believe it to be fair.

Theoretical Contributions

Our findings contribute to literatures on discrimination, meritocracy, and justice. First, we highlight understudied, yet foundational, principles of meritocracy: inputs must be controllable and relevant (e.g., Adams, 1965; Belmi et al., 2022; Son Hing et al., 2011). After extracting these latent criteria from existing theories of meritocracy, we demonstrate how they together shape perceptions of the fairness of using demographic criteria in selection—perceptions which can in turn foster discriminatory behavior. Our focus on these principles of meritocracy helps integrate work on stigma and attributions to discrimination, which has considered internality (e.g., controllability) of demographic attributes, with work from the justice literature, which has focused on face validity (e.g., relevance) of nondemographic attributes, such as interview questions and job tests. From this integration, we show that people assess discrimination by considering the selection procedure (*what criteria can be used?* procedural justice), rather than the outcome alone (*were groups treated equally?* distributive justice; cf. Gilliland, 1993).

Second, we shed light on an important reason why discrimination may persist: perceived controllability and relevance make some demographic attributes *feel* fair to select upon and nondiscriminatory. Moreover, we find that people weigh the perceived relevance of attributes more than they account for whether individuals can control the attributes in the first place (Studies 1a, 2a–2c, 4). Importantly, such perceived relevance remains a far cry from strict legal standards of bona-fide occupational qualifications; nevertheless, we find that individuals act upon these perceptions (see also Apfelbaum et al., 2017). In addition, we find that people express ambiguity about the fairness of mismatched attributes—those that are perceived to be high in relevance but low in controllability, or vice versa. In this way, the elements of “merit”—both controllability and relevance—may be played off one another, potentially leaving room for motivated reasoning and discriminatory decisions, nevertheless wrapped in a cloak of meritocracy.

Third, we provide useful descriptive evidence of individuals' perceptions of myriad demographic attributes, some of which are vastly understudied (e.g., religious beliefs, disability, caregiving status; Hebl et al., 2007). By assessing these attributes alongside well-studied ones (e.g., race, sex), we hope to have contributed to understanding how people perceive a wide variety of potentially discriminatory decisions. In particular, we find attributes vary in the degree to which people believe they *count* as discrimination. As such, we provide empirical evidence regarding when and why some forms of discrimination are considered relatively more fair and legitimate (Jetten, Iyer, et al., 2013).

Indeed, previous work suggests people may find social identity or demographic selection to be fair *unless* trained to think otherwise (Iyer et al., 2014). Our results suggest that meritocratic principles might be the terrain on which individuals negotiate and contest such legitimacy beliefs. For example, in the context of a novel attribute (Studies 3a–c), people's perceptions of fairness of using the attribute depended on experimentally manipulated perceptions of

controllability and relevance. This fits with the history of specific collective action movements. For instance, LGBTQ+ rights activists worked to generate support by countering a narrative of controllability and choice (Miceli, 2005; see also Haslam & Levy, 2006). Race- and sex-based civil rights movements have contested the idea that such attributes are relevant to performance.

In this way, our results may further suggest different mechanisms underlying the persistence of different types of discrimination. For instance, our results suggest that people find race and sex discrimination to be largely unfair, and yet such discrimination clearly persists in society (e.g., Bertrand & Mullainathan, 2004); that is, despite race and sex discrimination being perceived as unfair, implicit and explicit prejudice may lead individuals to subvert meritocracy and discriminate anyway. In contrast, our results suggest that age, disability, and caregiving discrimination may be driven not only by prejudice but also by perceptions of relevance suggesting these are indeed meritocratic attributes for selection; thus, no subversion is required. Identifying these underlying differences across legitimacy appraisals of attributes may help craft more tailored antidiscrimination interventions (Apfelbaum et al., 2016).

Practical Implications

Practically speaking, our results suggest that decision makers may place higher importance on relevance of demographic attributes than on controllability. As such, when a demographic attribute, such as caregiving status, is *perceived* as relevant to future performance, people may use information about that attribute to inform their decision, despite applicants' lack of control over such status—and despite legal protections. To the extent the relevance principle of meritocracy is prioritized, perhaps due to a business case mindset (Belmi & Pfeffer, 2015; Belmi & Schroeder, 2021) or perhaps because it is primary to meritocratic organization (with controllability reflecting a secondary overlay of individualist beliefs onto meritocracy), interventions might mitigate such prioritization by ensuring that demographic information is removed entirely from decision makers' materials.

In other cases, interventions might help decision makers assess relevance to an appropriate *legal* standard, rather than relying on individual perceptions of relevance alone. For instance, in Studies 3a–c, the described demographic attribute did not meet the standard of legal relevance to the job at hand, and yet we find that people's perceptions of relevance continue to affect their decisions. Indeed, relevance is foundational to U.S. discrimination law: permissible use of protected attributes involves extremely strict standards of relevance known as “bona-fide occupation qualifications.” Mere inability to perform a work function is not enough; employers are legally required to make accommodations for workers who are disabled, pregnant, or otherwise affected by their protected attribute status (e.g., religious observance; *Young v. United Parcel Service*, 2015). Further, in order to legally use such attributes in hiring, employers bear the burden of proving impact on a business *necessity*, the ability to perform a core function, rather than impact on mere competitiveness or success (*Wilson v. Southwest Airlines*, 1981). Again, meritocratic principles—that selection outcomes are based on relevant inputs—undergird such rationales. Training managers to understand this strict legal standard of relevance may be beneficial to reducing discrimination. Future work may also consider how people think about relevance of demographic

attributes in relationship to specific contexts and jobs, given this legal standard.

Here, we focused on hiring as an important context for considering discrimination and selection decisions and found that even those making hiring decisions in their jobs (relative experts; Study 2c) are influenced by relevance and controllability beliefs. Future work might consider whether the effects of controllability and relevance are even stronger outside of the more formalized (and scrutinized) domain of hiring; with less oversight and fewer guidelines, decision makers may be even more likely to rely on their own implicit assessments of merit, driven by perceived controllability and relevance. Ultimately, such assessments may perpetuate social inequity and further exclusion of members of nondominant demographic groups.

Limitations and Future Directions

The Origins of Relevance and Control Beliefs

The findings of the present work should be considered in light of its limitations. First, individuals' perceptions of controllability and relevance not only differ from others' perceptions but also differ from objective controllability and relevance. For instance, people believe that building a network or achieving an elite education is controllable, merely requiring time and effort from individuals (Grayson & Baldwin, 2011; Lin, 1999; Putnam, 1993). Further still, people believe that personal networks and elite education credentials signal fit and potential for high performance and are therefore relevant to the job at hand (Burks et al., 2015; Rivera, 2016; Williams et al., 1993). However, research also shows that both networks and elite educational opportunities are shaped largely by preexisting endowments (Castilla & Rissing, 2019; McDonald, 2011; Phillips et al., 2020; Reardon, 2018; Rivera, 2016), and that they do not objectively affect performance for most jobs (Collins, 2019; Shwed & Kalev, 2014). Similarly, people often *report* that they believe selecting employees who share demographic attributes with their current workforce (e.g., age, gender, race) will facilitate social integration (O'Reilly et al., 1989). However, such matching of demographic attributes of the workforce can *actually* be quite damaging to performance and other outcomes (Avery et al., 2008; Goldman et al., 2006; Phillips, 2014). The origins of these lay beliefs deserve further investigation (see also Lassetter et al., 2021).

A related limitation is that while we considered hiring expertise (Study 2c), *legal* expertise was not explicitly considered in the current work. Study 2b showed that regardless of the descriptive legality of a demographic attribute, controllability and relevance determined whether that attribute is fair to use in selection processes, and Study 3b extended these findings by explicitly informing participants that the target attribute was illegal to use. However, formal legal training may affect the extent to which people recognize and understand the legal status of attributes. For instance, we find that participants are more likely to treat some demographic attributes (e.g., race and sex) as clearly unfair to use, perhaps in part due to their salient legal status, while failing to recognize others which are nevertheless legally protected (e.g., age and caregiving status). As such, our work reveals that, *despite* legal protection, some demographic attributes are perceived as fair to use, due to perceived relevance and/or controllability. Further, some demographic

attributes are perceived to be unfair to use, despite the *lack* of legal protection (e.g., sexual orientation was unprotected during some of our data collection). Thus, our results suggest that what discrimination law says is fair and what people believe to be fair may not always align.

Our results also suggest that, beyond helping people learn which categories are protected by the law, interventions might focus on managing people's underlying perceptions of controllability and relevance. For instance, in U.S. discrimination law, protected categories are largely protected because individuals cannot easily control their ascribed membership in the category (e.g., race, sex). Indeed, the U.S. legal system is particularly concerned with avoiding decisions based on "accidents of birth" and immutability (Civil Rights Act, 1964), and modern discrimination law hinges on debates surrounding this very principle (e.g., Clarke, 2015; Green, 2017; Hoffman, 2011). As such, the logic of meritocracy appears: Only inputs that are individually controlled (and relevant, as described above) should merit reward or punishment.

The Role of Motivated Reasoning

In contrast with expertise, it is also likely that motivated reasoning plays a role in perceptions of controllability and relevance, and thus fairness. Ample prior work has demonstrated that people use the ideology of meritocracy to legitimate their own biased decisions, as well as systemic inequity (Augoustinos et al., 2005; Batruch et al., 2019; Castilla & Benard, 2010; Crandall & Eshleman, 2003; Knowles & Lowery, 2012; Ledgerwood et al., 2011; McCoy & Major, 2007; Phillips & Lowery, 2018, 2020; Savani & Rattan, 2012; Son Hing et al., 2011; Uhlmann & Cohen, 2005). Such motivated reasoning can even shape beliefs about specific criteria, including what people believe counts as "diverse" when they are told to seek diversity (Unzueta & Binning, 2012; Unzueta et al., 2012; see also Danbold & Unzueta, 2020). While we did not explore motivated antecedents directly, our results suggest another way in which meritocratic principles can perpetuate injustice: To the extent individuals are motivated to perceive demographic attributes as fair to use, they may turn to controllability and relevance to license their own preferences. Future work might consider how these specific procedural principles of meritocracy may enable motivated "casuistry," allowing individuals to launder their bias as legitimate (Norton et al., 2004).

At the same time, previous work on such "constructed criteria" has largely focused on how motivated reasoning shapes prioritization among attributes that—a priori—are widely agreed upon as appropriate criteria for assessing merit (language vs. math skills; Uhlmann & Cohen, 2005; see also Norton et al., 2004). That is, previous research has considered how people make *relative* decisions about the fairness of criteria that categorically reflect merit. Importantly, however, engaging in such legitimization practices—whether selfishly motivated or not—requires that people decide *which* attributes count as meritorious in the first place. We expand the set of attributes beyond *prima facie* legitimate criteria (e.g., math skills) to consider demographic attributes instead. Indeed, some perspectives on meritocracy suggest that demographic group or social identity criteria are *necessarily* antimeritocratic. Our results instead fit with Jetten, Iyer, et al.'s (2013) perspective, which suggests that demographic criteria are not seen as inherently anti-meritocratic; rather, such criteria are contested terrain. Indeed, our

results suggest that relevance and controllability may be the levers through which group categories become constructed instead as individual characteristics, and therefore indicators of merit.

Ultimately, future work might attempt to unpack the chicken/egg motivational question our results present. On the one hand, we find evidence that, for a novel attribute (i.e., about which individuals do not have existing motives), directly manipulated relevance and controllability affect perceptions of fairness and willingness to discriminate using that attribute (Studies 3a–c). In Studies 2a–c, we also find that perceptions of relevance and controllability relate to perceived fairness, above and beyond participants' own group membership or ideology. On the other hand, previous work has demonstrated that the desire to use an attribute (e.g., prejudice) itself can motivate justifications (Hegarty & Golden, 2008; Uhlmann & Cohen, 2005), and in our data, individuals certainly can differ regarding which attributes they believe are relevant and controllable (see also Online Supplemental Material). Moreover, the presence of socially accepted justifications itself can motivate the expression of prejudice (Crandall & Eshleman, 2003; see also Crandall et al., 2002). Our results suggest that meritocratic ideology in particular may provide the framework for such social acceptance, thereby licensing prejudice and discrimination such that these acts do not *feel* discriminatory.

In addition to prejudice, our work also speaks to the related but distinct phenomenon of essentialism (Haslam & Whelan, 2008; Prentice & Miller, 2007). In particular, part work has identified naturalness (belief that a category is biologically based and immutable) and inductive potential (the belief that category membership corresponds to some underlying essence shared by all members) as key dimensions of essentialism (Rothbart & Taylor, 1992; see also Haslam et al., 2000). Naturalness mirrors (un)controllability, while inductive potential may relate to relevance. At the same time, whereas inductive potential emphasizes the correspondence between group membership and an underlying essence, relevance as a principle of meritocracy emphasizes the correspondence between a selection criteria and the goal of selection (generally performance; Adams, 1965). Moreover, whereas naturalness and inductive potential are both positively associated with essentialism, their meritocracy corollaries are differently associated with fairness: negative for uncontrollability, but still positive for relevance. Overall, the analog between these core dimensions of essentialist beliefs and justice perceptions may suggest an even more fundamental role in human cognition. For instance, both principles may be involved in cognitions about social groups, with meritocracy in turn telling people how to *use* these dimensions (see also Prentice & Miller, 2007).

Finally, to consider the principles of meritocracy as a predictive structure for understanding perceptions of discrimination, we tested our hypotheses in a U.S. context. Relatedly, while we consider a wide range of demographic attributes, we do so in relative isolation. However, this prevents us from being able to generalize our results to other legal or cultural contexts, particularly those that do not endorse meritocracy beliefs. Future work might consider additional cultural contexts, with different ideological preferences and different attributes of interest (see also Salter et al., 2018; Savani & Rattan, 2012; Skurka et al., 2020). Future work might also consider multiple attributes simultaneously, to help understand the experiences of individuals who share multiple, intersecting attributes (e.g., a well-educated, younger, Black woman; Rosette, de Leon, et al., 2018).

Indeed, through cluster analysis (Studies 2a–2b), we find that some demographic attributes are perceived to be more similar to one another than others, but future work may focus on exploring the effects of sharing two or more demographic attributes within the same cluster, as well as between clusters.

Conclusion

Discrimination persists in part due to favoritism and animus toward demographic groups, which lead individuals to subvert fair, meritocratic principles in decision-making and day-to-day interactions (Bertrand & Mullainathan, 2004; Pager & Shepherd, 2008; Uhlmann & Cohen, 2005). However, even when meritocratic principles are upheld, discrimination may persist: people believe some kinds of demographic attributes to be meritocratic in the first place. Specifically, we find that people rely on core principles of meritocracy—controllability and relevance—to assess fairness of using demographic attributes. In turn, people can perceive even legally protected categories, such as age, disability, or caregiving responsibilities, as fair to use in selection. Thus, we find that because people believe that some demographic attributes meet the principles of meritocratic criteria, discrimination itself feels fair.

References

- Abrams, D., Swift, H. J., & Drury, L. (2016). Old and unemployable? How age-based stereotypes affect willingness to hire job candidates. *Journal of Social Issues, 72*(1), 105–121. <https://doi.org/10.1111/josi.12158>
- Adams, J. S. (1965). Inequity in social exchange. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 267–299). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60108-2](https://doi.org/10.1016/S0065-2601(08)60108-2)
- Apfelbaum, E. P., Grunberg, R., Halevy, N., & Kang, S. (2017). From ignorance to intolerance: Perceived intentionality of racial discrimination shapes preferences for colorblindness versus multiculturalism. *Journal of Experimental Social Psychology, 69*, 86–101. <https://doi.org/10.1016/j.jesp.2016.08.002>
- Apfelbaum, E. P., Stephens, N. M., & Reagans, R. E. (2016). Beyond one-size-fits-all: Tailoring diversity approaches to the representation of social groups. *Journal of Personality and Social Psychology, 111*(4), 547–566. <https://doi.org/10.1037/pspi0000071>
- Arrow, K. J. (1972). Models of job discrimination. In A. Pascal (Ed.), *Racial discrimination in economic life*. Lexington Books.
- Arvey, R. D., & Sackett, P. R. (1993). Fairness in selection: Current developments and perspectives. In N. Schmitt & W. Borman (Eds.), *Personnel selection* (pp. 171–202). Jossey-Bass.
- Augoustinos, M., Tuffin, K., & Every, D. (2005). New racism, meritocracy and individualism: Constraining affirmative action in education. *Discourse & Society, 16*(3), 315–340. <https://doi.org/10.1177/0957926505051168>
- Avery, D. R., McKay, P. F., & Wilson, D. C. (2008). What are the odds? How demographic similarity affects the prevalence of perceived employment discrimination. *Journal of Applied Psychology, 93*(2), 235–249. <https://doi.org/10.1037/0021-9010.93.2.235>
- Batruch, A., Autin, F., & Butera, F. (2019). The paradoxical role of meritocratic selection in the perpetuation of social inequalities at school. In J. Jetten & K. Peters (Eds.), *The social psychology of inequality* (pp. 123–137). Springer. https://doi.org/10.1007/978-3-030-28856-3_8
- Bauer, T. N., Truxillo, D. M., Sanchez, R. J., Craig, J. M., Ferrara, P., & Campion, M. A. (2001). Applicant reactions to selection: Development of the selection procedural justice scale (SPJS). *Personnel Psychology, 54*(2), 387–419. <https://doi.org/10.1111/j.1744-6570.2001.tb00097.x>
- Beauvois, J. E. L., & Dubois, N. (1988). The norm of internality in the explanation of psychological events. *European Journal of Social Psychology, 18*(4), 299–316. <https://doi.org/10.1002/ejsp.2420180402>

- Belmi, P., & Pfeffer, J. (2015). How “organization” can weaken the norm of reciprocity: The effects of attributions for favors and a calculative mindset. *Academy of Management Discoveries*, 1(1), 36–57. <https://doi.org/10.5465/amd.2014.0015>
- Belmi, P., Phillips, L., Laurin, K., & Engstrom, H. (2022). *Is it fair to get ahead by playing politics? The flexibility of merit enables self-serving judgments of political maneuvering* [Unpublished manuscript].
- Belmi, P., & Schroeder, J. (2021). Human “resources”? Objectification at work. *Journal of Personality and Social Psychology*, 120(2), 384–417. <https://doi.org/10.1037/pspi0000254>
- Berger, J., & Fişek, M. H. (2006). Diffuse status characteristics and the spread of status value: A formal theory. *American Journal of Sociology*, 111(4), 1038–1079. <https://doi.org/10.1086/498633>
- Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4), 991–1013. <https://doi.org/10.1257/0002828042002561>
- Brockner, J., Konovsky, M., Cooper-Schneider, R., Folger, R., Martin, C., & Bies, R. J. (1994). Interactive effects of procedural justice and outcome negativity on victims and survivors of job loss. *Academy of Management Journal*, 37(2), 397–409. <https://doi.org/10.5465/256835>
- Burks, S. V., Cowgill, B., Hoffman, M., & Housman, M. G. (2015). The value of hiring through employee referrals. *The Quarterly Journal of Economics*, 130(2), 805–839. <https://doi.org/10.1093/qje/qjv010>
- Bylsma, W. H., Major, B., & Cozzarelli, C. (1995). The influence of legitimacy appraisals on the determinants of entitlement beliefs. *Basic and Applied Social Psychology*, 17(1–2), 223–237. <https://doi.org/10.1080/01973533.1995.9646141>
- Castilla, E. J., & Benard, S. (2010). The paradox of meritocracy in organizations. *Administrative Science Quarterly*, 55(4), 543–676. <https://doi.org/10.2189/asqu.2010.55.4.543>
- Castilla, E. J., & Rissing, B. A. (2019). Best in class: The returns on application endorsements in higher education. *Administrative Science Quarterly*, 64(1), 230–270. <https://doi.org/10.1177/0001839218759965>
- Civil Rights Act of 1964, § 703 (Equal Employment Opportunity Commission 1964).
- Clarke, J. A. (2015). Against immutability. *The Yale Law Journal*, 125(1), 2–102. <https://www.yalelawjournal.org/article/against>
- Collins, R. (2019). *The credential society: An historical sociology of education and stratification*. Columbia University Press. <https://doi.org/10.7312/coll19234>
- Colquitt, J. A., & Rodell, J. B. (2015). Measuring justice and fairness. In R. S. Cropanzano & M. L. Ambrose (Eds.), *The Oxford handbook of justice in the workplace* (pp. 187–202). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199981410.013.8>
- Crandall, C. S. (1994). Prejudice against fat people: Ideology and self-interest. *Journal of Personality and Social Psychology*, 66(5), 882–894. <https://doi.org/10.1037/0022-3514.66.5.882>
- Crandall, C. S., & Eshleman, A. (2003). A justification-suppression model of the expression and experience of prejudice. *Psychological Bulletin*, 129(3), 414–446. <https://doi.org/10.1037/0033-2909.129.3.414>
- Crandall, C. S., Eshleman, A., & O’Brien, L. (2002). Social norms and the expression and suppression of prejudice: The struggle for internalization. *Journal of Personality and Social Psychology*, 82(3), 359–378. <https://doi.org/10.1037/0022-3514.82.3.359>
- Croon, M. A., & van Veldhoven, M. J. (2007). Predicting group-level outcome variables from variables measured at the individual level: A latent variable multilevel model. *Psychological Methods*, 12(1), 45–57. <https://doi.org/10.1037/1082-989X.12.1.45>
- Crosby, F. J. (2008). Sex discrimination at work. In J. C. Chrisler, C. Golden, & P. D. Rozee (Eds.), *Lectures on the psychology of women* (pp. 43–57). McGraw-Hill.
- Danbold, F., & Huo, Y. J. (2017). Men’s defense of their prototypicality undermines the success of women in STEM initiatives. *Journal of Experimental Social Psychology*, 72, 57–66. <https://doi.org/10.1016/j.jesp.2016.12.014>
- Danbold, F., & Unzueta, M. M. (2020). Drawing the diversity line: Numerical thresholds of diversity vary by group status. *Journal of Personality and Social Psychology*, 118(2), 283–306. <https://doi.org/10.1037/pspi0000182>
- Davey, L. M., Bobocel, D. R., Hing, L. S. S., & Zanna, M. P. (1999). Preference for the Merit Principle Scale: An individual difference measure of distributive justice preferences. *Social Justice Research*, 12(3), 223–240. <https://doi.org/10.1023/A:1022148418210>
- Deutsch, M. (1975). Equity, equality, and need: What determines which value will be used as the basis of distributive justice? *Journal of Social Issues*, 31(3), 137–149. <https://doi.org/10.1111/j.1540-4560.1975.tb01000.x>
- Devine, P. G., Forscher, P. S., Austin, A. J., & Cox, W. T. (2012). Long-term reduction in implicit race bias: A prejudice habit-breaking intervention. *Journal of Experimental Social Psychology*, 48(6), 1267–1278. <https://doi.org/10.1016/j.jesp.2012.06.003>
- Dipboye, R. L., & Colella, A. (Eds.). (2005). *The dilemmas of workplace discrimination at work: The psychological and organizational bases* (pp. 425–462). Psychology Press.
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, 57(6), 1082–1090. <https://doi.org/10.1037/0022-3514.57.6.1082>
- Eagly, A. H., & Heilman, M. E. (2016). Gender and leadership: Introduction to the special issue. *The Leadership Quarterly*, 27(3), 349–353. <https://doi.org/10.1016/j.leaqua.2016.04.002>
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383. <https://doi.org/10.2307/2666999>
- Feather, N. T. (1999). Judgments of deservingness: Studies in the psychology of justice and achievement. *Personality and Social Psychology Review*, 3(2), 86–107. https://doi.org/10.1207/s15327957pspr0302_1
- Feather, N. T. (2006). *Values, achievement, and justice: Studies in the psychology of deservingness*. Springer Science+Business Media.
- Finkelstein, L. M., Burke, M. J., & Raju, M. S. (1995). Age discrimination in simulated employment contexts: An integrative analysis. *Journal of Applied Psychology*, 80(6), 652–663. <https://doi.org/10.1037/0021-9010.80.6.652>
- Foladare, I. S. (1969). A clarification of “ascribed status” and “achieved status”. *The Sociological Quarterly*, 10(1), 53–61. <https://doi.org/10.1111/j.1533-8525.1969.tb02061.x>
- Ghumman, S., & Jackson, L. (2010). The downside of religious attire: The Muslim headscarf and expectations of obtaining employment. *Journal of Organizational Behavior*, 31(1), 4–23. <https://doi.org/10.1002/job.601>
- Gift, K., & Gift, T. (2015). Does politics influence hiring? Evidence from a randomized experiment. *Political Behavior*, 37(3), 653–675. <https://doi.org/10.1007/s11109-014-9286-0>
- Gilliland, S. W. (1993). The perceived fairness of selection systems: An organizational justice perspective. *Academy of Management Review*, 18(4), 694–734. <https://doi.org/10.2307/258595>
- Gilliland, S. W. (1995). Fairness from the applicant’s perspective: Reactions to employee selection procedures. *International Journal of Selection and Assessment*, 3(1), 11–18. <https://doi.org/10.1111/j.1468-2389.1995.tb00002.x>
- Goldman, B. M., Gutek, B. A., Stein, J. H., & Lewis, K. (2006). Employment discrimination in organizations: Antecedents and consequences. *Journal of Management*, 32(6), 786–830. <https://doi.org/10.1177/0149206306293544>
- Grayson, C., & Baldwin, D. (2011). *Leadership networking: Connect, collaborate, create*. Wiley.
- Green, M. W., Jr. (2017). Same-Sex sex and immutable traits: Why Obergefell v. Hodges clears a path to protecting gay and lesbian employees

- from workplace discrimination under Title VII. *The Journal of Gender, Race & Justice*, 20(1), 1–52. https://engagedscholarship.csuohio.edu/fac_articles/888
- Hansen, R. S. (1992). A multidimensional scale for measuring business ethics: A purification and refinement. *Journal of Business Ethics*, 11(7), 523–534. <https://doi.org/10.1007/BF00881445>
- Haslam, N., & Levy, S. R. (2006). Essentialist beliefs about homosexuality: Structure and implications for prejudice. *Personality and Social Psychology Bulletin*, 32(4), 471–485. <https://doi.org/10.1177/0146167205276516>
- Haslam, N., Rothschild, L., & Ernst, D. (2000). Essentialist beliefs about social categories. *British Journal of Social Psychology*, 39(1), 113–127. <https://doi.org/10.1348/014466600164363>
- Haslam, N., & Whelan, J. (2008). Human natures: Psychological essentialism in thinking about differences between people. *Social and Personality Psychology Compass*, 2(3), 1297–1312. <https://doi.org/10.1111/j.1751-9004.2008.00112.x>
- Hebl, M. R., King, E. B., Glick, P., Singletary, S. L., & Kazama, S. (2007). Hostile and benevolent reactions toward pregnant women: Complementary interpersonal punishments and rewards that maintain traditional roles. *Journal of Applied Psychology*, 92(6), 1499–1511. <https://doi.org/10.1037/0021-9010.92.6.1499>
- Hegarty, P., & Golden, A. M. (2008). Attributional beliefs about the controllability of stigmatized traits: Antecedents or justifications of prejudice? *Journal of Applied Social Psychology*, 38(4), 1023–1044. <https://doi.org/10.1111/j.1559-1816.2008.00337.x>
- Heilman, M. E. (1983). Sex bias in work settings: The lack of fit model. *Research in Organizational Behavior*, 5, 269–298. <https://psycnet.apa.org/record/1984-10927-001>
- Hochschild, J. L. (1996). *Facing up to the American dream: Race, class, and the soul of the nation*. Princeton University Press. <https://doi.org/10.1515/9781400821730>
- Hoffman, S. (2011). *The importance of immutability in employment discrimination law*. Case Western Reserve University School of Law.
- Iyer, A., Jetten, J., Branscombe, N. R., Jackson, S., & Youngberg, C. (2014). The difficulty of recognizing less obvious forms of group-based discrimination. *Group Processes & Intergroup Relations*, 17(5), 577–589. <https://doi.org/10.1177/1368430214522139>
- Jetten, J., Branscombe, N. R., Iyer, A., & Asai, N. (2013). Appraising gender discrimination as legitimate or illegitimate: Antecedents and consequences. In M. K. Ryan & N. R. Branscombe (Eds.), *Handbook of gender and psychology* (pp. 306–322). SAGE.
- Jetten, J., Iyer, A., Branscombe, N. R., & Zhang, A. (2013). How the disadvantaged appraise group-based exclusion: The path from legitimacy to illegitimacy. *European Review of Social Psychology*, 24(1), 194–224. <https://doi.org/10.1080/10463283.2013.840977>
- Jetten, J., Schmitt, M. T., Branscombe, N. R., Garza, A. A., & Mewse, A. J. (2011). Group commitment in the face of discrimination: The role of legitimacy appraisals. *European Journal of Social Psychology*, 41(1), 116–126. <https://doi.org/10.1002/ejsp.743>
- Jones, R. G. (Ed.). (2013). *Nepotism in organizations*. Routledge. <https://doi.org/10.4324/9780203805886>
- Kang, S. K., DeCelles, K. A., Tilcsik, A., & Jun, S. (2016). Whiteness résumés: Race and self-presentation in the labor market. *Administrative Science Quarterly*, 61(3), 469–502. <https://doi.org/10.1177/0001839216639577>
- Kay, A. C., Whitson, J. A., Gaucher, D., & Galinsky, A. D. (2009). Compensatory control: Achieving order through the mind, our institutions, and the heavens. *Current Directions in Psychological Science*, 18(5), 264–268. <https://doi.org/10.1111/j.1467-8721.2009.01649.x>
- Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (Vol. 15, pp. 192–240). University of Nebraska Press.
- Kluegel, J. R., & Smith, E. R. (1986). *Beliefs about inequality: Americans' views of what is and what ought to be*. Aldine de Gruyter.
- Knowles, E. D., & Lowery, B. S. (2012). Meritocracy, self-concerns, and Whites' denial of racial inequity. *Self and Identity*, 11(2), 202–222. <https://doi.org/10.1080/15298868.2010.542015>
- Kouchaki, M., Smith-Crowe, K., Brief, A., & Sousa, C. (2013). Seeing green: Mere exposure to money triggers a business decision frame and unethical outcomes. *Organizational Behavior and Human Decision Processes*, 121(1), 53–61. <https://doi.org/10.1016/j.obhdp.2012.12.002>
- Krefting, L. A., & Brief, A. P. (1976). The impact of applicant disability on evaluative judgments in the selection process. *Academy of Management Journal*, 19(4), 675–680. <https://doi.org/10.5465/255802>
- Lasseter, B., Hehman, E., & Neel, R. (2021). The relevance appraisal matrix: Evaluating others' relevance. *Journal of Personality and Social Psychology. Advance online publication*. <https://doi.org/10.1037/pspi0000359>
- Ledgerwood, A., Mandisodza, A. N., Jost, J. T., & Pohl, M. J. (2011). Working for the system: Motivated defense of meritocratic beliefs. *Social Cognition*, 29(3), 322–340. <https://doi.org/10.1521/soco.2011.29.3.322>
- Lee, M., Pitesa, M., Pillutla, M. M., & Thau, S. (2018). Perceived entitlement causes discrimination against attractive job candidates in the domain of relatively less desirable jobs. *Journal of Personality and Social Psychology*, 114(3), 422–442. <https://doi.org/10.1037/pspi0000114>
- Lerner, M. J. (1977). The justice motive: Some hypotheses as to its origins and forms. *Journal of Personality*, 45(1), 1–52. <https://doi.org/10.1111/j.1467-6494.1977.tb00591.x>
- Lerner, M. J. (1980). *The belief in a just world: A fundamental delusion*. Plenum.
- Lin, N. (1999). Social networks and status attainment. *Annual Review of Sociology*, 25, 467–487. <https://doi.org/10.1146/annurev.soc.25.1.467>
- Lupfer, M. B., Weeks, K. P., Doan, K. A., & Houston, D. A. (2000). Folk conceptions of fairness and unfairness. *European Journal of Social Psychology*, 30(3), 405–428. [https://doi.org/10.1002/\(SICI\)1099-0992\(200005/06\)30:3<405::AID-EJSP997>3.0.CO;2-U](https://doi.org/10.1002/(SICI)1099-0992(200005/06)30:3<405::AID-EJSP997>3.0.CO;2-U)
- Major, B. (1994). From social inequality to personal entitlement: The role of social comparisons, legitimacy appraisals, and group membership. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 26, pp. 293–355). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60156-2](https://doi.org/10.1016/S0065-2601(08)60156-2)
- Major, B., & Dover, T. L. (2016). Attributions to discrimination: Antecedents and consequences. In T. D. Nelson (Ed.), *Handbook of prejudice, stereotyping, and discrimination* (pp. 213–239). Psychology Press.
- Major, B., Gramzow, R. H., McCoy, S. K., Levin, S., Schmader, T., & Sidanius, J. (2002). Perceiving personal discrimination: The role of group status and legitimizing ideology. *Journal of Personality and Social Psychology*, 82(3), 269–282. <https://doi.org/10.1037/0022-3514.82.3.269>
- Major, B., Quinton, W. J., & McCoy, S. K. (2002). Antecedents and consequences of attributions to discrimination: Theoretical and empirical advances. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 34, pp. 251–330). Academic Press. [https://doi.org/10.1016/S0065-2601\(02\)80007-7](https://doi.org/10.1016/S0065-2601(02)80007-7)
- Major, B., Quinton, W. J., & Schmader, T. (2003). Attributions to discrimination and self-esteem: Impact of group identification and situational ambiguity. *Journal of Experimental Social Psychology*, 39(3), 220–231. [https://doi.org/10.1016/S0022-1031\(02\)00547-4](https://doi.org/10.1016/S0022-1031(02)00547-4)
- Martí, M. W., Bobier, D. M., & Baron, R. S. (2000). Right before our eyes: The failure to recognize non-prototypical forms of prejudice. *Group Processes & Intergroup Relations*, 3(4), 403–418. <https://doi.org/10.1177/136843020003004005>
- McCoy, S. K., & Major, B. (2007). Priming meritocracy and the psychological justification of inequality. *Journal of Experimental Social Psychology*, 43(3), 341–351. <https://doi.org/10.1016/j.jesp.2006.04.009>
- McDonald, S. (2011). What's in the “old boys” network? Accessing social capital in gendered and racialized networks. *Social Networks*, 33(4), 317–330. <https://doi.org/10.1016/j.socnet.2011.10.002>
- McNamara, D. S., & McDaniel, M. A. (2004). Suppressing irrelevant information: Knowledge activation or inhibition? *Journal of Experimental*

- Psychology: Learning, Memory, and Cognition*, 30(2), 465–482. <https://doi.org/10.1037/0278-7393.30.2.465>
- McNamee, S. J., & Miller, R. K. (2009). *The meritocracy myth*. Rowman & Littlefield.
- Miceli, M. (2005). *Standing out, standing together: The social and political impact of gay-straight alliances*. Routledge.
- Miller, F. D. (2001). *Moral knowledge*. Cambridge University Press.
- Nadler, J. T., & Kufahl, K. M. (2014). Marital status, gender, and sexual orientation: Implications for employment hiring decisions. *Psychology of Sexual Orientation and Gender Diversity*, 1(3), 270–278. <https://doi.org/10.1037/sgd0000050>
- Noon, M. (2007). The fatal flaws of diversity and the business case for ethnic minorities. *Work, Employment and Society*, 21(4), 773–784. <https://doi.org/10.1177/0950017007082886>
- North, M. S. (2019). A GATE to understanding “older” workers: Generation, age, tenure, experience. *The Academy of Management Annals*, 13(2), 414–443. <https://doi.org/10.5465/annals.2017.0125>
- Norton, M. I., Vandello, J. A., & Darley, J. M. (2004). Casuistry and social category bias. *Journal of Personality and Social Psychology*, 87(6), 817–831. <https://doi.org/10.1037/0022-3514.87.6.817>
- O’Reilly, C. A., III, Caldwell, D. F., & Barnett, W. P. (1989). Work group demography, social integration, and turnover. *Administrative Science Quarterly*, 34(1), 21–37. <https://doi.org/10.2307/2392984>
- Pachankis, J. E., Hatzenbuehler, M. L., Wang, K., Burton, C. L., Crawford, F. W., Phelan, J. C., & Link, B. G. (2018). The burden of stigma on health and well-being: A taxonomy of concealment, course, disruptiveness, aesthetics, origin, and peril across 93 stigmas. *Personality and Social Psychology Bulletin*, 44(4), 451–474. <https://doi.org/10.1177/0146167217741313>
- 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>
- Phillips, K. W. (2014). How diversity works. *Scientific American*, 311(4), 42–47. <https://doi.org/10.1038/scientificamerican1014-42>
- Phillips, L. T., & Jun, S. (2021). Why benefiting from discrimination is less recognized as discrimination. *Journal of Personality and Social Psychology*. Advance online publication. <https://doi.org/10.1037/pspi0000298>
- Phillips, L. T., & Lowery, B. S. (2018). Herd invisibility: The psychology of racial privilege. *Current Directions in Psychological Science*, 27(3), 156–162. <https://doi.org/10.1177/0963721417753600>
- Phillips, L. T., & Lowery, B. S. (2020). I ain’t no fortunate one: On the motivated denial of class privilege. *Journal of Personality and Social Psychology*, 119(6), 1403–1422. <https://doi.org/10.1037/pspi0000240>
- Phillips, L. T., Stephens, N. M., Townsend, S. S. M., & Goudeau, S. (2020). Access is not enough: Cultural mismatch persists to limit first-generation students’ opportunities for achievement throughout college. *Journal of Personality and Social Psychology*, 119(5), 1112–1131. <https://doi.org/10.1037/pspi0000234>
- Pitesa, M., & Pillutla, M. M. (2019). Socioeconomic mobility and talent utilization of workers from poorer backgrounds: The overlooked importance of within-organization dynamics. *The Academy of Management Annals*, 13(2), 737–769. <https://doi.org/10.5465/annals.2017.0115>
- Prentice, D. A., & Miller, D. T. (2007). Psychological essentialism of human categories. *Current Directions in Psychological Science*, 16(4), 202–206. <https://doi.org/10.1111/j.1467-8721.2007.00504.x>
- Purdie-Vaughns, V., Steele, C. M., Davies, P. G., Dittmann, R., & Crosby, J. R. (2008). Social identity contingencies: How diversity cues signal threat or safety for African Americans in mainstream institutions. *Journal of Personality and Social Psychology*, 94(4), 615–630. <https://doi.org/10.1037/0022-3514.94.4.615>
- Putnam, R. (1993). The prosperous community: Social capital and public life. *The American Prospect*, 13(4), 35–42. <http://www.prospect.org/print/vol13>
- Quinn, D. M., & Crocker, J. (1999). When ideology hurts: Effects of belief in the protestant ethic and feeling overweight on the psychological well-being of women. *Journal of Personality and Social Psychology*, 77(2), 402–414. <https://doi.org/10.1037/0022-3514.77.2.402>
- Reardon, S. F. (2018). *The widening academic achievement gap between the rich and the poor* (pp. 536–550). Routledge. <https://doi.org/10.4324/9780429499821-33>
- Reidenbach, R. E., & Robin, D. P. (1990). Toward the development of a multidimensional scale for improving evaluations of business ethics. *Journal of Business Ethics*, 9(8), 639–653. <https://doi.org/10.1007/BF00383391>
- Rhode, D. L. (2010). *The beauty bias: The injustice of appearance in life and law*. Oxford University Press.
- Rivera, L. A. (2016). *Pedigree: How elite students get elite jobs*. Princeton University Press.
- Rodin, M. J., Price, J. M., Bryson, J. B., & Sanchez, F. J. (1990). Asymmetry in prejudice attribution. *Journal of Experimental Social Psychology*, 26(6), 481–504. [https://doi.org/10.1016/0022-1031\(90\)90052-N](https://doi.org/10.1016/0022-1031(90)90052-N)
- Rosette, A. S., Akinola, M., & Ma, A. (2018). Subtle discrimination in the workplace: Individual-level factors and processes. In A. J. Colella & E. B. King (Eds.), *Oxford library of psychology. The Oxford handbook of workplace discrimination* (pp. 7–24). Oxford University Press.
- Rosette, A. S., de Leon, R. P., Koval, C. Z., & Harrison, D. A. (2018). Intersectionality: Connecting experiences of gender with race at work. *Research in Organizational Behavior*, 38(6033), 1–22. <https://doi.org/10.1016/j.riob.2018.12.002>
- Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173–220). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60357-3](https://doi.org/10.1016/S0065-2601(08)60357-3)
- Rothbart, M., & Taylor, M. (1992). Category labels and social reality: Do we view social categories as natural kinds? In G. R. Semin & K. Fiedler (Eds.), *Language, interaction and social cognition* (pp. 11–36). Sage Publications.
- Salter, P. S., Adams, G., & Perez, M. J. (2018). Racism in the structure of everyday worlds: A cultural-psychological perspective. *Current Directions in Psychological Science*, 27(3), 150–155. <https://doi.org/10.1177/0963721417724239>
- Salter, P. S., Hirsch, K. A., Schlegel, R. J., & Thai, L. T. (2016). Who needs individual responsibility? Audience race and message content influence third-party evaluations of political messages. *Social Psychological & Personality Science*, 7(1), 29–36. <https://doi.org/10.1177/1948550615590447>
- Sanchez, D. T., Chaney, K. E., Manuel, S. K., Wilton, L. S., & Remedios, J. D. (2017). Stigma by prejudice transfer: Racism threatens White women and sexism threatens men of color. *Psychological Science*, 28(4), 445–461. <https://doi.org/10.1177/0956797616686218>
- Sanchez, J. I., & Brock, P. (1996). Outcomes of perceived discrimination among Hispanic employees: Is diversity management a luxury or a necessity? *Academy of Management Journal*, 39(3), 704–719. <https://doi.org/10.5465/2566660>
- Savani, K., & Rattan, A. (2012). A choice mind-set increases the acceptance and maintenance of wealth inequality. *Psychological Science*, 23(7), 796–804. <https://doi.org/10.1177/0956797611434540>
- Schmitt, M. T., Branscombe, N. R., Postmes, T., & Garcia, A. (2014). The consequences of perceived discrimination for psychological well-being: A meta-analytic review. *Psychological Bulletin*, 140(4), 921–948. <https://doi.org/10.1037/a0035754>
- Seidel, M. D. L., Polzer, J. T., & Stewart, K. J. (2000). Friends in high places: The effects of social networks on discrimination in salary negotiations. *Administrative Science Quarterly*, 45(1), 1–24. <https://doi.org/10.2307/2666977>
- Shwed, U., & Kalev, A. (2014). Are referrals more productive or more likeable? Social networks and the evaluation of merit. *American*

- Behavioral Scientist*, 58(2), 288–308. <https://doi.org/10.1177/0002764213503330>
- Skurka, C., Winett, L. B., Jarman-Miller, H., & Niederdeppe, J. (2020). All things being equal: Distinguishing proportionality and equity in moral reasoning. *Social Psychological & Personality Science*, 11(3), 374–387. <https://doi.org/10.1177/1948550619862261>
- Smither, J. W., Reilly, R. R., Millsap, R. E., Pearlman, K., & Stoffey, R. W. (1993). Applicant reactions to selection procedures. *Personnel Psychology*, 46(1), 49–76. <https://doi.org/10.1111/j.1744-6570.1993.tb00867.x>
- Son Hing, L. S., Bobocel, D. R., Zanna, M. P., Garcia, D. M., Gee, S. S., & Oraziotti, K. (2011). The merit of meritocracy. *Journal of Personality and Social Psychology*, 101(3), 433–450. <https://doi.org/10.1037/a0024618>
- Spector, P. E. (1994). *Job satisfaction survey*. University of South Florida.
- Thibaut, J. W., & Walker, L. (1975). *Procedural justice: A psychological analysis*. Lawrence Erlbaum.
- Tonidandel, S., & LeBreton, J. M. (2011). Relative importance analysis: A useful supplement to regression analysis. *Journal of Business and Psychology*, 26(1), 1–9. <https://doi.org/10.1007/s10869-010-9204-3>
- Tyler, T. R. (2006). Psychological perspectives on legitimacy and legitimation. *Annual Review of Psychology*, 57, 375–400. <https://doi.org/10.1146/annurev.psych.57.102904.190038>
- Uhlmann, E., & Cohen, G. L. (2005). Constructed criteria: Redefining merit to justify discrimination. *Psychological Science*, 16(6), 474–480. <https://doi.org/10.1111/j.0956-7976.2005.01559.x>
- 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>
- Unzueta, M. M., & Binning, K. R. (2012). Diversity is in the eye of the beholder: How concern for the in-group affects perceptions of racial diversity. *Personality and Social Psychology Bulletin*, 38(1), 26–38. <https://doi.org/10.1177/0146167211418528>
- Unzueta, M. M., Knowles, E. D., & Ho, G. C. (2012). Diversity is what you want it to be: How social-dominance motives affect construals of diversity. *Psychological Science*, 23(3), 303–309. <https://doi.org/10.1177/0956797611426727>
- van den Bos, K., & Lind, E. A. (2002). Uncertainty management by means of fairness judgments. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 34, pp. 1–60). Academic Press. [https://doi.org/10.1016/S0065-2601\(02\)80003-X](https://doi.org/10.1016/S0065-2601(02)80003-X)
- Walster, E., Berscheid, E., & Walster, G. W. (1973). New directions in equity research. *Journal of Personality and Social Psychology*, 25(2), 151–176. <https://doi.org/10.1037/h0033967>
- Wang, L., Zhong, C. B., & Murnighan, J. K. (2014). The social and ethical consequences of a calculative mindset. *Organizational Behavior and Human Decision Processes*, 125(1), 39–49. <https://doi.org/10.1016/j.obhdp.2014.05.004>
- Wax, A. L. (1999). Discrimination as accident. *Indiana Law Journal*, 74, 1129–1214. http://scholarship.law.upenn.edu/faculty_scholarship/747
- Weiner, B. (1995). *Judgments of responsibility: A foundation for a theory of social conduct*. Guilford Press.
- Williams, C. R., Labig, C. E., & Stone, T. H. (1993). Recruitment sources and posthire outcomes for job applicants and new hires: A test of two hypotheses. *Journal of Applied Psychology*, 78(2), 163–172. <https://doi.org/10.1037/0021-9010.78.2.163>
- Wilson v. Southwest Airlines Co., 517 F. Supp. 292 (N.D. Tex. 1981).
- Young, M. D. (1958). *The rise of the meritocracy*. Thames and Hudson.
- Young v. United Parcel Service, 575 U.S. (2015)
- Zhong, C. B. (2011). The ethical dangers of deliberative decision making. *Administrative Science Quarterly*, 56(1), 1–25. <https://doi.org/10.2189/asqu.2011.56.1.001>

Received May 24, 2021

Revision received October 20, 2021

Accepted November 3, 2021 ■