

THE UNIVERSITY OF CHICAGO

FAIRNESS AND MERIT IN THE SELECTION OF ADVANTAGED AND
DISADVANTAGED APPLICANTS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY

DAVID MAURICIO MUNGUIA GOMEZ

CHICAGO, ILLINOIS

JUNE 2023

Copyright © 2023

David Mauricio Munguia Gomez

All rights reserved

Table of Contents

| | |
|----------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| List of Tables | iv |
| List of Figures | v |
| Acknowledgments..... | vi |
| Overview..... | 1 |
| Chapter 1: The Policy-People Gap: Decision-makers choose policies that favor different applicants than they select when making individual decisions..... | 4 |
| Chapter 2: How Interpretations of Socioeconomic Advantage and Disadvantage Influence College Admissions | 54 |
| Chapter 3: Luck that Builds Merit | 71 |
| Appendix A: Tables | 86 |
| Appendix B: Figures..... | 91 |
| References..... | 107 |
| Note on Supplementary Materials | 119 |

List of Tables

| | |
|-------------------------------------------------------------------------|----|
| Table 1.1 Correlations between Independent and Dependent Variables..... | 86 |
| Table 1.2 Overview of Main and Supplementary Studies..... | 87 |
| Table 3.1 Overview of Studies..... | 88 |
| Table 3.2 Means and Standard Deviations of Dependent Variables..... | 89 |
| Table 3.3 Correlations between Dependent Variables..... | 90 |

List of Figures

| | |
|---------------------------------------------------------------------------------|-----|
| Figure 1.1 Theoretical Model for Policy Decisions and Individual Decisions..... | 91 |
| Figure 1.2 Stimuli (Study 1)..... | 92 |
| Figure 1.3 Stimuli (Study 4)..... | 93 |
| Figure 1.4 Choice by Decision Type and Intervention Condition..... | 94 |
| Figure 1.5 Forest Plot of Internal Meta-Analysis on Choice..... | 95 |
| Figure 2.1 Experimental Stimuli..... | 96 |
| Figure 2.2 Results of Meta-Analysis..... | 97 |
| Figure 2.3 Effort Adjustment by Political Attitudes..... | 98 |
| Figure 2.4 Adjustment Directions..... | 99 |
| Figure 2.5 Income Distributions..... | 100 |
| Figure 3.1 Fairness Judgment..... | 101 |
| Figure 3.2 Choice to Invite..... | 102 |
| Figure 3.3 Effort..... | 103 |
| Figure 3.4 Skill..... | 104 |
| Figure 3.5 Internal..... | 105 |
| Figure 3.6 Tournament Prediction..... | 106 |

Acknowledgments

I am lucky in countless ways that have allowed me to write this dissertation and pursue an academic career.

I am lucky to have David Munguía and Rosa Reyna Gómez as my parents. Their commitment to my education, including driving me across the U.S.-Mexico border every day, made it possible for me to even read and write in English. I hope to one day fully appreciate all the sacrifices they have made for me.

I am lucky to have had Eloim Munguía as my grandfather. His hard work and commitment to his family changed the course of my life before I was even born. I hope to be as dedicated and humble as he was.

I am lucky to have had Emma Levine as my adviser. Her expertise and guidance made it possible for me to do research that I am proud of and secure a job that will allow me to continue doing so. Her friendship helped me push through challenging moments in my life and the uncertainties of this career. Her unconditional support and kindness helped me stay productive, motivated, and happy. I aspire to have the same kind of impact on others that Emma has had on me.

I am lucky to have had Jane Risen, Chris Hsee, and Taylor Phillips guide my work. They are joyful, kind, and thoughtful. Their feedback and encouragement greatly improved this dissertation. I strive to emulate them as researchers and mentors.

I am lucky to have had scholars early on who gave me opportunities and helped me develop as a researcher – Bethany Burum, Dan Gilbert, Eldar Shafir, Mike North, and Susan Fiske. Like them, I hope to inspire and support others on their academic journeys.

I am lucky to have trained and conducted my research surrounded by so many scholars that were generous with their time and thoughts – Andrew Meyer, Anuj Shah, Ayelet Fishbach, Berkeley Dietvorst, Chris Bryan, Dan Bartels, Erika Kirgios, George Wu, James Dungan, Margaret Echelbarger, Nick Epley, Oleg Urminsky, Reid Hastie, and Thomas Talhelm. They improved my thinking, confidence, and research.

I am lucky to have so many caring friends with whom to share this journey –Adam Mastroianni, Akshina Banerjee, Cintia Hinojosa, Dani Goya-Tocchetto, Diag Davenport, Erika Kirgios, Kyle Dobson, Kristina Wald, Rafael Batista, Sami Kassirer, Xilin Li, and Yuji Winet. They make this career easier and more fun.

I am lucky to have Lucie Wright as my life partner. She is my greatest cheerleader and support. I hope to help her achieve her goals and aspirations as she helps me fulfill mine.

Overview

Who do we admit to college? Hire for a job and promote? Give an award, loan, or even visa to? These are some of the most consequential decisions we regularly make about others. What determines their outcome?

Candidates' merit – their perceived achievements, qualifications, ability, and effort – largely predicts who gets these rewards and opportunities. For example, applicants with higher standardized test scores are more likely to be admitted to college, particularly to more selective institutions, and job candidates with more years of experience are more likely to receive callbacks for an interview (Bertrand & Mullainathan, 2004). However, these decisions are not only a function of candidates' merit.

Candidates' circumstances also play a role. By circumstances I mean aspects of a candidate's environment that are advantageous or disadvantageous, that is, associated with better or worse outcomes and experiences. Examples of circumstances that we know are consequential are family income, geographic location, primary schooling, and family stability (Chetty, Hendren, Kline, & Saez, 2014). The effect of these circumstances on decisions about candidates may appear straightforward: disadvantaged candidates may be less likely than advantaged candidates to be selected. Indeed, research on, for example, discrimination, stereotyping, and cultural mismatch support this notion (see, e.g., Kraus, Torrez, Park, & Ghayebi, 2019; Phillips, Stephens, Townsend, & Goudeau, 2020; Rivera & Tilcsik, 2016). However, candidates' circumstances may influence decisions in others ways that benefit not only the advantaged, but also the disadvantaged.

My dissertation seeks to unearth and understand the ways in which candidates' circumstances affect the outcomes of decisions about who to reward and give opportunity to. It suggests that circumstances affect these decisions in two ways – by influencing what people find fair and the attributions they draw about candidates.

Chapter 1 discovers one determinant of whether people favor advantaged or disadvantaged candidates in their decisions – whether they are choosing between policies or between specific individuals. In this work, which focuses on college admissions and workplace hiring, I find that decision-makers are more likely to favor disadvantaged applicants over advantaged applicants with objectively higher achievements when choosing between selection policies than choosing between individual applicants. I find that this gap is driven in part by shifting standards of fairness across the two types of decisions. When choosing between individuals, compared to choosing between policies, decision-makers are more likely to prioritize what is fair to individuals (microjustice) over what is fair in the aggregate (macrojustice).

Chapter 2 examines how people evaluate a college applicant before versus after learning about their socioeconomic circumstances. This work sheds light on people's beliefs about candidates' socioeconomic circumstances and how people interpret candidates' merit in light of their advantages and disadvantages. I find that people adjust their decisions and impressions asymmetrically – they positively adjust their decision to admit a disadvantaged applicant more than they negatively adjust their decision to admit an advantaged applicant. This difference appears to stem from inferences about the applicants' effort, which are influenced by the mental models they hold about how socioeconomic (dis)advantage affects applicants.

Chapter 3 examines how people make decisions in situations in which candidates' circumstances gave them the opportunity to develop merit. People generally want to select people for rewards and opportunities based on their merit and not on their luck (Almås, Cappelen, Sørensen, & Tungodden, 2010; Starmans, Sheskin, Bloom, 2017). However, achievements can stem from merit that was made possible by luck. When faced with candidates whose achievements stem from luck made possible by merit, how do people choose to reward them? How does this compare to situations in which candidates' achievements reflect just luck or just ability? This chapter begins to answer these questions.

Chapter 1: The Policy-People Gap: Decision-makers choose policies that favor different applicants than they select when making individual decisions

Abstract

This work documents a contemporary organizational problem—a gap between selection policies and individual selection decisions—and suggests one intervention to address it. In college admissions and workplace hiring contexts, we find that decision-makers are more likely to favor disadvantaged applicants over applicants with objectively higher achievements when choosing between selection policies than choosing between individual applicants. We document this policy–people gap among admissions officers, working professionals, and lay people using both within-subject and between-subject designs and across a range of stimuli. We find that the gap is driven in part by shifting standards of fairness across the two types of decisions. When choosing between individuals, compared to choosing between policies, decision-makers are more likely to prioritize what is fair to individuals (a microjustice standard of fairness) over what is fair in the aggregate (a macrojustice standard of fairness). As a result, an intervention that has decision-makers prioritize the same standard of fairness across the decisions mitigates the policy–people gap. This research helps us understand why decision-makers’ choices so frequently violate espoused organizational policies and suggests one way to increase the representation of disadvantaged groups in organizations. The supplementary materials linked at the end include additional analyses for the studies reported here as well as additional studies.

Introduction

In the mid-1990s, half of Americans favored affirmative action policies for women and racial minorities (NBC News & The Wall Street Journal, 1995, 1997). Yet, in that same time period, only 34% of Americans actually thought that it was a good idea to hire a woman over a man if they were both equally qualified, and only 20% thought that a college should deliberately favor a Black applicant over an equally qualified White applicant (CBS News & New York Times, 1997). Why was there a gap between respondents' preferred policies and whom they would favor in specific cases? In the present research, we suggest that such a gap existed, and persists,¹ because decision-makers prioritize different standards of fairness when making decisions about selection policies versus decisions about which specific individuals to select. Our goals in the present paper are threefold. First, we seek to establish that decisions about selection policies diverge from decisions about which specific individuals to select. We tested this hypothesis in the contexts of college admissions and workplace hiring and found robust evidence that the decisions diverge, even when these two types of decisions have identical consequences. Across 16 preregistered studies with 10,883 participants, we found that admissions officers, employees with hiring experience in the technology industry, and members of the general public are more likely to choose a policy that favors lower-scoring, disadvantaged applicants over higher-scoring, advantaged applicants than they are to select a specific lower-scoring, disadvantaged applicant over a higher-scoring, advantaged applicant. We refer to this systematic divergence as the “policy–people gap.”

Our second goal is to explain why a gap between selection policies and individual selection decisions exists. We theorize that when people make individual selection decisions they

¹ Comparison of recent polls about affirmative action: <https://bit.ly/3fNq273>

prioritize a “microjustice” standard of fairness, that is, what is fair to the individuals under consideration. Microjustice reflects concerns about whether opportunities (e.g., admissions spots, jobs) are allocated in proportion to individuals’ specific qualities (e.g., their accomplishments and productivity). However, the relatively more abstract nature of policy decisions leads people to consider a different standard of fairness. When people make policy decisions, we theorize that they become less motivated by a microjustice standard and become more motivated by a “macrojustice” standard of fairness, that is, what is fair in the aggregate. Macrojustice reflects concerns about how opportunities ought to be distributed overall, for example, whether those opportunities are distributed equally between groups or whether some minimum for a group is reached. These standards of fairness lead decision-makers to favor different applicants in their policy and individual decisions.

Our third goal is to present an intervention that aligns individual decisions with policy decisions by leveraging our theoretical account. Our account suggests that when decision-makers rely on the same standards of fairness across decisions, the gap between the decisions will be reduced. We crafted an intervention in which decision-makers learned about both microjustice and macrojustice and were told that their individual decisions should help accomplish their institution’s priority of macrojustice. This minimal intervention led significantly more decision-makers to make individual decisions consistent with a policy favoring lower-scoring, disadvantaged applicants, thereby helping bridge the policy–people gap.

In pursuing these goals, we make several theoretical and practical contributions. First, we document a robust and persistent gap between the selection decisions people make when choosing policies and when choosing individual applicants. This finding advances our understanding of why individuals and organizations often fail to live up to their espoused moral

standards: Decision-makers care about multiple standards of fairness, but the macrojustice standard of fairness that is relevant to achieving equality across groups is often deprioritized when making individual selection decisions.

Our findings dovetail with interdisciplinary research that has examined divergences between decision-makers' endorsement of broader principles and the concrete implementation of those principles (e.g., Bell & Hartmann, 2007; Dixon, Durrheim, & Tredoux, 2007; Smith & Mayorga-Gallo, 2017; Yogeewaran & Dasgupta, 2014). Some existing work has found, for example, that people support principles of racial equality but oppose specific actions to achieve it (e.g., Dixon et al., 2007). Some have argued that this gap stems from the ambiguity afforded by endorsing abstract principles or a failure to consider what implementing those principles actually implies (Bell & Hartmann, 2007; Smith & Mayorga-Gallo, 2017). Another stream of research on "want–should" conflicts has suggested that people fail to implement their abstract principles because in contextualized, affect-rich contexts, they focus on what they want to do rather than what they should do (Bazerman & Tenbrunsel, 2011; Milkman, Rogers, & Bazerman, 2008). Thus far, however, the mechanisms underlying such gaps have been "hotly debated" (Tuch & Hughes, 1996: 724), with scholars noting that "there is not yet a consensus about [their] nature and causes" (Dixon et al., 2007: 868).

We provide a new perspective on why principle– implementation gaps exist and how to mitigate them. We establish that decision-makers' choices diverge even when the implementation of organizational policies is clear and equivalent to the individual decision in question. In doing so, we suggest that gaps between the endorsement of principles and their specific implementation cannot be solely explained by previous accounts, such as a failure to think through implementation (Bell & Hartmann, 2007; Smith & Mayorga-Gallo, 2017) or being tempted by

what one wants to do (Milkman et al., 2008). Further, we provide a novel theoretical account to explain why policy and individual decisions diverge even when decision-makers think carefully about what they should do when making both types of decisions. We theorize and find evidence that people focus on what they should do (i.e., what they find fair to do) across both policy and individual decisions but that the relevant standard of fairness shifts across them.

This work also sheds light on the psychology of policymaking. There is a desire among behavioral scientists and organizational behavior scholars to influence policy as evidenced, for example, by the creation of the peer-reviewed journal, *Behavioral Science & Policy*, the growing number of nudge units around the world (Benartzi et al., 2017), and the interdisciplinary publications that make policy recommendations (e.g., Rattan, Savani, Chugh, & Dweck, 2015; Soll, Keeney, & Larrick, 2013), investigate policymaking (e.g., Hideg, Michela, & Ferris, 2011; Milkman, Mazza, Shu, Tsay, & Bazerman, 2012; Rogers & Bazerman, 2008), and examine the relationship between research and policy (e.g., Adler & Jermier, 2005; Pearce & Huang, 2012; Tannenbaum, Fox, & Rogers, 2017). However, little research to date has actually examined the psychological factors that facilitate or impede policies from having an impact on the decisions that they intend to guide. The present research contributes to this body of work by elucidating how policymaking is distinct from other types of decisions: when people make policy decisions, they are more motivated by a macrojustice standard of fairness, which can be at odds with their motivations when making one-off decisions.

Practically, this work also helps identify psychological drivers of organizational and societal inequality. We examined our research questions in the contexts of college admissions and workplace hiring, which represent critical moments of economic, social, and labor market stratification (Arce & Segura, 2015; Chetty, Friedman, Saez, Turner, & Yagan, 2017; Stephens,

Markus, & Phillips, 2014). Thus, we are able to link the judgments and decision-making of organizational gatekeepers with broader patterns of organizational and societal inequality. We devoted special attention to college admissions, which, while valuable in its own right, also helps explain the underrepresentation of women and racial minorities in the workplace. Practices and phenomena that are typically studied in the workplace also exist in institutions of higher learning, such as affirmative action policies (Hideg & Ferris, 2017), cultural mismatch (Phillips, Stephens, Townsend, & Goudeau, 2020), and cultural matching (Rivera, 2011). Moreover, some of the forces that shape underrepresentation and inequality in the workplace exist before applicants even apply to an organization (see, e.g., Cheryan, Ziegler, Montoya, & Jiang, 2017; Milkman, Akinola, & Chugh, 2015; Rivera & Stevens, 2013).

Theory

Selecting individuals for an organization entails evaluating them and deciding whom to select, whereas choosing a selection policy entails representing “goals, intentions, or statements” (Van Meter & Van Horn, 1975: 446) in order to make many individual selection decisions. We propose that this distinct function of policy decisions—to serve as a statement that guides future selection decisions—leads decision-makers to prioritize different standards of fairness when making policy versus individual decisions. We draw on construal-level theory (Trope & Liberman, 2010) and the concepts of microjustice and macrojustice (Brickman, Folger, Goode, & Schul, 1981) to explain the differences between the two decisions. Broadly, we posit that policy decisions are represented more abstractly than individual decisions, thus shifting decision-makers’ standard of fairness from what is fair to specific individuals being evaluated (microjustice) to what is a fair distribution in the aggregate (macrojustice). These shifting

standards of fairness influence subsequent selection decisions. We depict this theoretical account in Figure 1.1.

People can mentally represent objects, behaviors, and other people at different levels of construal (Trope & Liberman, 2003; 2010). A “high-level construal” is a representation that is relatively abstract, decontextualized, and focused on “central, superordinate, and goal-relevant features” (Soderberg, Callahan, Kochersberger, Amit, & Ledgerwood, 2015: 2). In contrast, a “low-level construal” is a concrete, contextualized, and detailed representation. Consider, for example, the act of moving into a new apartment. When construing this action at a high level, people described it as “starting a new life,” but, when construing it at a low level, they described it as “packing and carrying boxes” (Liberman & Trope, 1998). “Starting a new life” reflects the broader goals associated with moving, whereas “packing and carrying boxes” reflects the concrete actions involved in it. A number of dispositional and contextual features, such as a person’s mood (Gasper & Clore, 2002), power (Smith & Trope, 2006), and current social role (e.g., mother, manager [Reyt & Wiesenfeld, 2015]), can shift the degree to which people construe concepts at different levels (for reviews of construal, see Trope & Liberman, 2010; Wiesenfeld, Reyt, Brockner, & Trope, 2017).

Drawing on this fundamental distinction in people’s mental representations, we theorize that selection decisions focused on individuals versus policies differ in their level of construal. Both types of decisions reflect ways to expand an organization. However, policies are crafted to represent the broader goals of organizational expansion, whereas individual selection decisions reflect the specific concrete actions taken in pursuit of those goals. Individual selection decisions require evaluating specific applicants and determining their suitability for a particular opportunity or job. These decisions, therefore, operate at a relatively lower level of construal.

We theorize that people rely on a microjustice standard of fairness when making decisions that are construed at a low level. Microjustice reflects “principles of correspondence” and thus demands that an appropriate relation be established between an individual’s qualities or “inputs” (e.g., accomplishments, productivity, and need) and their reward or “outcome” (e.g., college admission, job placement; Brickman et al., 1981: 178). As such, microjustice is individuating—it requires knowing individuals’ qualities in order to determine what they ought to get. This means that microjustice is only possible in decisions and contexts that attend to specific information and, as attention is drawn to that information, it may seem to be the appropriate standard of fairness. In the context of selection decisions, which seek to determine applicants’ suitability for a job or opportunity, inputs pertaining to perceived merit (e.g., ability, accomplishments, and experience) are focal (see, e.g., Bohnet, Van Geen, & Bazerman, 2016; Quadlin, 2018), and thus prioritizing microjustice entails focusing on signals of merit.²

In contrast to individual selection decisions, policies guide behavior across time, situations, and people. They seek to, for example, guide behavior today or in 10 years, in the context of hiring for an executive position or an entry-level one, by decision-makers with distinct experiences and knowledge. For policies to represent the values and goals of policy-makers in a way that will traverse these variable factors, we theorize that they are necessarily represented at a higher level of construal than individual decisions.

We theorize that the relatively more abstract nature of policies prompts decision-makers to rely on a different standard of fairness. Specifically, we theorize that when thinking at a higher

² The relationship between microjustice and merit may be specific to the selection context. In decision contexts in which people seek to foster personal development or social relations, rather than maximize economic productivity, other inputs, like those pertaining to need, may be more relevant (see, e.g., Deutsch, 1975; Wagstaff, Huggins, & Perfect, 1993).

level of construal, decision-makers are less motivated by what is fair to specific individuals (microjustice) and instead are more motivated by what is fair in the aggregate (e.g., in society, their organization, in an industry). In other words, that decision-makers prioritize macrojustice, which is concerned with what the distribution of resources ought to be like overall, such as no one having less than a certain amount (a “minimum principle”) or having an appropriate balance in the allocation of resources between groups (a “balance principle”; Brickman et al., 1981: 190). In contrast to microjustice, macrojustice does not attend to individuals’ specific attributes because what the overall distribution of resources ought to be like does not depend on those attributes.

The link between construal level and prioritizing microjustice versus macrojustice has not yet been established, but it is consistent with existing research examining allocations to individuals and groups (Chang, Kirgios, Rai, & Milkman, 2020; Colby, DeWitt, & Chapman, 2015; Ubel & Loewenstein, 1996). For example, when considering which specific patients should receive organs, decision-makers allocate organs to the patients with the highest likelihood of survival; this result reflects a correspondence between the likelihood of survival (an input) and receiving an organ (the outcome), consistent with a microjustice standard of fairness (Colby et al., 2015; Ubel & Loewenstein, 1996). However, when considering groups of patients, decision-makers allocate organs more equally across the groups; this result reflects a balance principle, consistent with a macrojustice standard of fairness (Colby et al., 2015; Ubel & Loewenstein, 1996). The broader group-based decision led decision-makers to prioritize a macrojustice standard of fairness that was not as compelling at the individual level, namely, that everyone should have the same chance at receiving an organ.

The link between construal level and prioritizing macrojustice versus microjustice is also consistent with research on the relationship between construal level and values (Luguri, Napier, & Dovidio, 2012; Mahfud, Badea, Verkuyten, & Reynolds, 2018; Mentovich, Yudkin, Tyler, & Trope, 2016; Napier & Luguri, 2013; Yogeewaran & Dasgupta, 2014). When thinking more abstractly, people across the political spectrum increase their endorsement of values related to equality (Luguri et al., 2012; Napier & Luguri, 2013), which reflect the notion of macrojustice. Similarly, when thinking more abstractly, people make ethical choices about individuals that rely less on their specific traits, such as granting the same level of speech protections to U.S. citizens and non-citizens or making similar salary cuts for local and foreign workers (Mentovich et al., 2016), thus reflecting lower endorsement of microjustice.

We further theorize that the relative prioritization of microjustice versus macrojustice within individual versus policy decisions influences which applicants decision-makers favor. When faced with a choice between an advantaged applicant with higher objective achievements and a disadvantaged applicant with lower objective achievements, the more decision-makers prioritize microjustice relative to macrojustice, the more likely they will be to favor the applicant with higher objective achievements.³ Microjustice is associated with selecting the higher-achieving applicant because it attends to signals of the individuals' merit to determine whether they ought to be selected, whereas macrojustice is associated with selecting the disadvantaged

³ By achievement, we mean objective metrics, such as scores on the SAT or on a coding assignment, that signal an applicants' qualities, like their college readiness or ability to perform a job. Though these metrics are based on unambiguously right and wrong answers (and thus objective), that does not necessarily make them accurate measures of applicants' qualities.

applicant because it is concerned with balance and equality in the aggregate. Overall, our theoretical account leads us to make the following predictions:

***Hypothesis 1:** Decision-makers favor higher-achieving, advantaged individuals over lower-achieving, disadvantaged individuals more often when making a decision about individuals than when making a decision about policies.*

***Hypothesis 2:** A greater relative prioritization of microjustice (versus macrojustice) accounts for the relationship between making an individual versus policy decision and the likelihood of favoring higher-achieving, advantaged individuals over lower-achieving, disadvantaged individuals.*

In Hypothesis 2, we posit that the relative prioritization of microjustice versus macrojustice changes between the decisions and influences choice. It need not be the case that decision-makers are more concerned about macrojustice than microjustice when choosing policies or that they are more concerned about microjustice than macrojustice when choosing individuals—only that decision-makers differentially prioritize macrojustice and microjustice between the decisions. This hypothesis assumes that most decision-makers are concerned about both microjustice and macrojustice to some extent but that making a policy versus an individual decision changes their relative prioritization. This assumption is consistent with existing work showing that people value multiple, and sometimes conflicting, standards of fairness (Colquitt, Conlon, Wesson, Porter, & Ng, 2001) and that the use of these standards can change as a function of contextual factors, such as organizational structure (Ambrose & Schminke, 2003), hunger (Huppert, Shaw, & Decety, 2020), and social relationships (Deutsch, 1975; Rai & Fiske, 2011).

We also consider how to align policy and individual decisions. We theorize that the gap in choice between policy and individual decisions stems from the different standards of fairness that underlie them. Therefore, we predict that aligning decision-makers to use the same standard of fairness across their decisions should attenuate the gap in choice. We propose that:

***Hypothesis 3:** Decision-makers favor higher-achieving, advantaged individuals over lower-achieving, disadvantaged individuals in their individual and policy decisions at more similar rates when they rely on the same standard of fairness (microjustice or macrojustice) across the decisions than when they do not.*

Paper Overview

Across four main studies ($N = 2,698$; all preregistered) we tested these hypotheses. In Studies 1 and 2, we found support for Hypothesis 1, with samples of college admissions officers and employees in the technology industry. In Study 3, we measured microjustice and macrojustice concerns and found support for Hypothesis 2, using mediation analyses. In Study 4, we tested a macrojustice intervention and found support for Hypotheses 2 and 3.

In our Supplementary Findings section, which follows Study 4, we highlight the results from our supplementary studies (15 studies; $n = 8,477$) that address alternative explanations and the robustness of our effects. Our four main studies provide the most direct tests of our hypotheses. However, we report our supplementary studies for transparency (we have no “file-drawer” of preregistered studies) and for the additional insights they provide into the policy–people gap. We highlight takeaways from our supplementary studies in the Supplementary Findings section, and we list each study’s key findings in Table 1.2. Readers interested in further details about any of our studies can consult the supplementary materials (SM). We have also

made all of our data, analysis files, preregistrations, and survey materials available online (<https://bit.ly/3tJB1q7>).

Experimental Design

Across our studies, participants made a choice either between two specific individuals or between two policies that would favor those individuals. Within those choices, we set up a tension between prioritizing achievement and taking into account applicants' circumstances; the latter can allow one to realize values of equality and diversity or to assess merit based on circumstances. Hence, participants choosing between individuals faced a choice between a higher-achieving, advantaged individual and a lower-achieving, disadvantaged individual, and participants choosing between policies faced a choice between a policy that would favor higher-achieving, advantaged individuals and a policy that would favor lower-achieving, disadvantaged individuals. We surveyed 35 college admissions officers (see SM 2.15) to gauge the frequency with which actual decision-makers experience this tension and to ensure that we developed reasonable stimuli to capture it. Fifty percent of the admissions officers confirmed that they had previously faced a college admissions decision that pitted favoring applicants with greater qualifications who were relatively advantaged against applicants with lower qualifications who were relatively disadvantaged.

We also ran a recall study with 108 management faculty (see SM 2.14). The vast majority of faculty who completed our survey reported recent involvement in hiring (90% were involved within the last three years). We asked them to recall a time when they faced a trade-off between “candidates with higher objective metrics of success (e.g., many publications, excellent methodological training) or candidates with marginally lower metrics of success who may have been disadvantaged in some way (e.g., they lacked research funding, had a limited stipend,

belong to a group or social category that faces more obstacles in academia).” Eighty-six percent of the management faculty could recall making faculty hiring decisions that involved this trade-off, which is the same one we represent in our studies. Based on these two pilot studies, we posit that our operationalization reasonably reflects decisions and tensions that admissions officers and hiring parties actually face.

We made three important design choices to test our hypotheses as effectively and cleanly as possible. First, we restricted the decision to just two individuals or two policies and described those options based on easily quantifiable information. Without a doubt, in actual selection contexts, there is more than one way to select an option (e.g., ranking, from a pool, in isolation), and there is often more and richer information available about the options. However, by creating a relatively bare decision environment, we were able to cleanly test the causal effect of making a decision about individuals versus policies. Second, we held constant the information presented in the individual and policy decisions. Specifically, we used the same information to describe the individuals and the policies; we describe this procedure in detail in the study methods. Third, we sampled participant populations according to the goals of our studies. In Studies 1 and 2, we sampled decision-makers within organizations to establish the gap between policy and individual decisions and its generalizability. In Studies 3 and 4, as well as in the supplementary studies, we sought to make more granular comparisons, so we sampled large numbers of participants from Amazon’s Mechanical Turk ([mturk.com](https://www.mturk.com)) and Prolific Academic (prolific.co)—two online platforms that allow for high-quality, experimental data collection (Buhrmester, Talaifar, & Gosling, 2018; Peer, Brandimarte, Samat, & Acquisti, 2017). The participants in these platforms provide useful data for investigating varied organizational phenomena (e.g., hiring policies, as in Hideg & Ferris, 2017).

For studies with online participants, we implemented a variety of checks to ensure that participants were attentive and read all the instructions. The checks varied across the studies based on the approach we considered best at the time, but we preregistered all exclusions based on the checks to preclude “p-hacking.” Nevertheless, where relevant, we illustrate key analyses with all possible participants included (see Footnotes 5 and 7), which reveal that the results did not meaningfully change based on the specific exclusions.

Study 1: College Admissions

In Study 1, we tested whether college admissions officers are more likely to select lower-achieving, disadvantaged applicants over higher-achieving, advantaged applicants when selecting admissions policies than they are when making individual admissions decisions. We tested this hypothesis with a sample of admissions officers using both a within-subjects and between-subjects design. This mixed design allowed us to examine whether the same decision-maker would make divergent choices. We preregistered this study on AsPredicted.org (https://aspredicted.org/DQG_UDX).

Methods

Participants. We aimed to recruit as many participants as possible during a two-day period at a national conference for college admissions. We targeted conference attendees wearing lanyards that identified them as university or college admissions officers or as high school college counselors. We ended up with a final sample of 190 participants; 159 participants completed the survey in full, and 31 participants only completed parts of it (overall demographics: 59% female, $M_{\text{age}} = 37.25$, $SD = 9.83$). Nearly all of the participants (96.5%) had at least one year of experience in an admissions department ($M_{\text{AdmissionsExperience}} = 10.03$ years, $SD = 8.42$).

Procedure and materials. We distributed paper surveys to participants. We asked them to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. Then, participants made an admissions decision between two individuals followed by a decision between two admissions policies, or they made a decision between two admissions policies followed by an admissions decision between two individuals. That is, all participants made a decision about admissions policies and about individuals to admit, but the order of the decisions varied. This design allowed us to examine the hypothesis between subjects by comparing the initial choices across participants and within subjects by comparing participants' own first and second choices.

Before making their individual or policy decision, participants examined a table of information about the two applicants or two admissions policies. Participants saw the same information for both types of decisions. However, the table of information for the individual decision framed the information as describing, "two specific applicants for admission, Applicant A and Applicant B," whereas the table for the policy decision framed the information as describing the "types of applicants that are typically admitted under Policy A (Applicant Type A) and Policy B (Applicant Type B)." We depict the tables of information in Figure 1.2.

The applicants were identical in all respects except for their SAT scores and estimated household incomes. One applicant (or type of applicant) had higher scores and household income, while the other had lower scores and household income. Hence, participants chose between a higher-scoring, economically advantaged option and a lower-scoring, economically disadvantaged option. The specific numbers provided in the table (e.g., SAT and household income) were the same for all participants and were always associated with the same (type of

applicant (e.g., Applicant A always had an SAT score of 1200 and a household income of \$60,000).

After examining the table of information, participants indicated which of the two options they would select (“If you had to choose one [applicant to admit/admissions policy], which [applicant/policy] would you choose?”: A or B) and how likely they would be to choose either option on a scale of 0 (much more likely to [admit Applicant A / choose Policy A]) to 10 (much more likely to [admit Applicant B / choose Policy B]). We also asked participants to indicate which option would be the fairer one to select (“Which [applicant/policy] is it fairer to select?”: A or B). Measuring fairness provides initial insight into our account; we hypothesized that participants would make choices they believe are fair in both decision contexts but that they would use different standards of fairness to make those judgments.

After participants answered the questions for their first (individual or policy) decision, they proceeded to their second decision. They reviewed the same table of information, now framed for the other type of decision, and answered the same set of questions. We also collected a series of exploratory mechanism items (see SM 1.1). Finally, participants reported their demographic information and provided general details about their admissions experience and academic institutions.

Results

Some participants did not provide answers to all questions or their answers were not decipherable by research assistants transcribing the surveys. In order to maintain a higher level of statistical power, we considered all available responses for each analysis. Therefore, we have a different number of participants across analyses. For each analysis, we examined participants’ decisions in two ways. We compared (a) only participants’ first policy and individual choices (a

between-subjects comparison) and (b) participants' first policy and individual choices to their own subsequent choices in the other condition (a within-subjects comparison).⁴

Choice. Comparing only participants' first choices (between subjects), we found that participants chose the higher-scoring, higher-income option in the individual condition at a higher rate (70.11%) than in the policy condition (36.56%; $\chi^2 [1, 180] = 20.31, p < .001$). Comparing participants' own choices (within subjects), we also found this pattern of results (individual: 66.08% vs. policy 39.18%; McNemar test, $\chi^2 [1, 171] = 31.12, p < .001$).

Likelihood of choice. Between subjects, participants reported a greater likelihood of choosing the higher-scoring, higher-income option in the individual condition ($M = 5.97, SD = 2.22$) than in the policy condition ($M = 5.20, SD = 2.84; t[183] = 2.04, p = .043$). We also found this pattern of results within subjects ($M_{Individual} = 6.19, SD = 2.38$ vs. $M_{Policy} = 5.11, SD = 2.70; t[179] = 5.54, p < .001$).

Fair choice. Between subjects, participants indicated that the higher-scoring, higher-income option was the fair choice to select in the individual condition at a higher rate (53.75%) than in the policy condition (35.16%; $\chi^2 [1, 171] = 5.97, p = .015$). We also found this pattern of results within subjects (individual: 54.60% vs. policy: 35.58%; McNemar test, $\chi^2 [1, 163] = 18.84, p < .001$).

Discussion

We found that the majority of admissions officers preferred to admit higher-scoring, advantaged applicants over lower-scoring, disadvantaged applicants. However, when choosing

⁴ As indicated in our preregistration, we also tested for order effects. We did not find any differences in choice based on whether the individual or policy decision came first. We report these tests and their results in SM 1.1.

an admission policy, most officers preferred the policy that would admit lower-scoring, disadvantaged applicants over higher-scoring, advantaged applicants.

The admissions officers also indicated that admitting the higher-scoring, advantaged applicant was the fairer applicant to select more often than they indicated that the policy that would favor that applicant was the fairer policy. Hence, this study provides evidence that what decision-makers choose and what they think is fair to choose differs when selecting individual applicants versus policies. This finding is consistent with our theorizing.

Strikingly, we found that the policy–people gap persists both between subjects and within subjects. Forty percent of admissions officers made a subsequent choice about specific individuals (or policies) that diverged from their initial choice about policies (or specific individuals). This result is even more striking because the admissions officers made their second decision immediately after the first one, and the information used to describe each type of decision was identical. These results suggest that the gap cannot be easily explained by decision-makers attending to or having different information in each decision or by failing to consider how the policies would inform individual decisions. If those mechanisms drove the policy–people gap, we would not expect to see divergent choices when the same decision-makers make both decisions and the consequences of each were salient. Rather, these results are more consistent with the possibility that decision-makers are aware of their divergent choices and see them as reasonable because they believe different standards are relevant for each decision. We test this directly in Studies 3 and 4.

Study 2: Workplace Hiring

In Study 2, we tested Hypothesis 1 in a new selection context: hiring decisions within a technology start-up. Employees in the technology industry were randomly assigned to make

individual or policy decisions about whom to interview for a job. As in the admissions context of Study 1, the choice between job candidates presented a trade-off between prioritizing achievement and taking circumstances into account. We preregistered this study on AsPredicted.org (https://aspredicted.org/EGG_NRT).

Methods

Participants. We recruited United States-based and United Kingdom-based online participants on Prolific Academic for a small monetary sum. Using Prolific’s filters, we sought to recruit participants with experience working in technology-related fields. The majority of participants (68.2%) reported working in STEM and technology-related occupations, and about one-third (34.7%) reported experience in hiring ($M = 5.78$ years, $SD = 6.28$). We presented an attention check before assigning participants to a condition; those who failed the check were not allowed to continue with the study. Therefore, we are not able to conduct analyses with and without participants who failed the check. Our final sample included 360 participants (36.67% female, $M_{age} = 35.24$, $SD = 10.24$).

Procedure and materials. We asked participants to imagine that they were the head of human resources at a technology start-up and that they had to decide which job candidates to interview for their start-up’s software engineer position. Then, some participants chose one of two job candidates to interview, while other participants chose one of two interviewing policies.

Before making their decision, participants examined a table of information with the two options they could choose between. Participants saw the same table of information for both types of decisions, but the framing of the options differed slightly for each. For the individual decision, the table described “two candidates,” whereas for the policy decision, the table described the

“types of candidates that would be interviewed under each policy.” We provide the table of information in SM 1.3.

The two options described in the table were identical in all respects except for their “coding assignment score” and the “candidate gender.” Before seeing the table, participants read that technology companies often give a coding assignment to job candidates “to make sure that they can do the job they are applying for” and that their start-up “gives an assignment scored from 0 to 10, where 5 (the average) indicates the candidate is competent enough to do the job well.” One of the options had a male candidate with a score of 8.1, whereas the other option had a female candidate with a score of 8.0. We chose these stimuli because we anticipated that participants would face a tension between prioritizing achievement (by favoring the male candidate with an objectively higher score) and taking the applicants’ gender into account (by favoring the female candidate, aware that there are more barriers for women in STEM and that they are underrepresented in the industry). We counterbalanced whether the higher-scoring, male option or the lower-scoring, female option was labeled “A” and appeared on the left.

After examining the table of information, participants indicated which of the two options they would select (“If you had to choose one [candidate to interview/interviewing policy], which [candidate/policy] would you choose?": A or B) and which option would be the fairer one to select (“Which [candidate/interviewing policy] is it fairer to [interview/ select]?": A or B). We also collected a series of exploratory mechanism items (see SM 1.2). At the end of the study, participants reported demographic information (age, gender, and political attitudes), their current employment sector or industry, and years of hiring experience (if any).

Results

Choice. Participants chose the higher-scoring, male option in the individual condition at a higher rate (64.48%) than in the policy condition (52.54%; $\chi^2 [1, 360] = 5.29, p = .021$).

Fair choice. Participants indicated that the higher-scoring, male option was the fair option to choose in the individual condition at a higher rate (79.23%) than in the policy condition (56.50%; $\chi^2 [1, 360] = 21.40, p < .001$).

Discussion

We found that technology employees were more likely to interview—and believed it was fairer to interview—higher-scoring, male candidates, over slightly lower-scoring, female candidates when making decisions about individual candidates than when choosing an interviewing policy. These results provide further evidence of a gap between policy decisions and decisions about individuals (Hypothesis 1) and suggest that these results extend to the workplace. Moreover, these results are consistent with our theoretical account that decision-makers' choices reflect what they consider to be fair but that the standard of fairness differs across the decisions.

Study 3: Mediation by Shifting Standards of Fairness

In Study 3, we directly tested the mechanism proposed by our theoretical account, specifically, whether the policy–people gap is explained by decision-makers being relatively more motivated by microjustice than macrojustice when making individual versus policy decisions (Hypothesis 2). We also compared the individual and policy decisions in a more controlled way than in Studies 1 and 2. Specifically, we compared an individual admissions decision (choosing one of two applicants) to a policy decision that would only affect two applicants, thereby making the consequences of the decisions identical. Finding a policy–people gap in this more controlled paradigm would suggest that any differences between the decisions in

this study were due to how decision-makers think about policies rather than underlying features of typical policy decisions, such as the greater number of people they affect or the reduced identifiability of those people (Ritov & Zamir, 2014). We preregistered this study on AsPredicted.org (https://aspredicted.org/9K8_HXJ).

Methods

Participants. We recruited United States-based online participants on Amazon Mechanical Turk for a small monetary sum. We estimated requiring a sample size of 800 in order to run the intended mediation analyses. Based on Study S7, we assumed a small X-to-M effect and a large M-to-Y effect, which would require approximately 400 participants (Fritz & MacKinnon, 2007). However, because we also introduced another factor (randomizing the order in which the dependent variables and mechanism measures were collected), we doubled our target sample size to 800.

At the beginning of the study, we tested whether participants were able to see a sample table of information (we tested for this because we displayed the information about applicants and policies in a table). Participants who gave a response that did not match the words on the table were prevented from taking the rest of the survey. As indicated in our preregistration, our analyses excluded participants who failed any of the three following attention checks. Two checks pertained to the individual or policy decision that participants made. One check asked participants what their choice was about: “who to admit to college” (this was the correct answer based on the subject matter of the entire scenario), “... interview for a job,” “... nominate for a prize,” “... recommend for a promotion,” or “... choose for a sports team.” The other check asked which of two answer options participants chose between in the preceding screen: “two applicants” or “two policies.” The final check was embedded in the eight justice questions we

presented to participants and stated, “Please select the middle answer option (4) to show you are paying attention.” Our final sample comprised 802 participants ($M_{age} = 41.00$, $SD = 12.72$; 56.36% female).⁵

Procedure and materials. We randomly assigned participants to one of two conditions (decision type: individual or policy) in a between-subjects design. As in Study 1, we asked participants to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. We presented them with the same table of information as in Study 1 and counter-balanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option was labeled “A” and appeared on the left.

However, unlike in Study 1, participants in the individual and policy conditions made functionally equivalent decisions. The only thing that differed was the framing of the decision—whether participants conceptualized their choice as one about policies or individuals. Participants in the policy condition learned that they would choose between two admissions policies that “only affect one final pair of applicants this year.” Because participants might find it odd that the policies would only affect two applicants, we explained that the policy would “only affect one final pair of applicants” because “it is almost the end of admissions season.” We did not want

⁵ For robustness, we also conducted our analyses including the 115 participants (13% of the sample) who failed these checks. The results do not meaningfully change when we include these participants. To allow readers to compare to the results reported ahead, we report two key analyses here that include these 115 participants. Participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (56.49%) than participants choosing between policies (49.01%; $\chi^2 [1, 917] = 5.15$, $p = .023$). Participants indicated that their decision was (or would be) motivated by microjustice considerations significantly more when choosing between individuals ($M = 5.77$, $SD = 0.96$) than when choosing between equivalent policies ($M = 5.54$, $SD = 1.23$; $t[915] = 3.19$, $p = .002$). Conversely, they indicated that their decision was motivated (or would be) by macrojustice considerations marginally less when choosing between individuals ($M = 3.94$, $SD = 1.91$) than when choosing between equivalent policies ($M = 4.16$, $SD = 1.89$; $t[915] = 1.74$, $p = .082$).

participants to speculate whether the policies would affect more applicants down the line, so we further explained that they and their “fellow admissions officers can create a new policy next year.” Then, participants saw a table of information describing “the two specific applicants that would be admitted under each policy.” Participants in the individual condition similarly learned that “it is almost the end of admissions season” and that they were making a decision about “one final pair of applicants,” and then they proceeded to see the table of information describing the two specific applicants. After reading about each applicant or policy, participants responded to our dependent measures.

Choice and fairness. Participants selected one of the two applicants or policies (forced choice: A or B), indicated which applicant or policy was the fairer one to select (forced choice: A or B), and reported how likely they were to pick either option on a scale of 1 (much more likely to choose A) to 7 (much more likely to choose B).

Microjustice and macrojustice. Participants also reported how much their choice of applicant or policy was motivated by eight concerns related to macrojustice and microjustice on a scale of 1 (not at all) to 7 (extremely). We presented all eight items together in random order and counterbalanced whether they came before or after the choice, fairness, and likelihood questions. When the items came before making a choice, we asked participants how much their choice “would be motivated” by these concerns, whereas when the items came after making a choice, we asked how much the choice “was motivated” by the concerns.

The macrojustice items measured the extent to which participants were motivated by “making sure that there is a diverse population within my university,” “creating greater equality in the opportunities disadvantaged and advantaged applicants receive,” “balancing the number of opportunities given to disadvantaged and advantaged applicants,” and “making sure that there

are not too few disadvantaged applicants in my university,” whereas the microjustice items measured the extent to which participants were motivated by “selecting applicants based on their hard work,” “selecting applicants based on their qualifications,” “selecting applicants based on their potential to succeed,” and “selecting applicants based on their accomplishments.” We developed these items based on the original conceptualization of the macrojustice and microjustice constructs (Brickman et al., 1981) and our other studies (Studies S7–S9).

We conducted a confirmatory factor analysis to test the factor structure of the items. We assessed model fit with χ^2 , root-mean-square error of approximation (RMSEA), comparative fit index (CFI), and Tucker–Lewis index (TLI). The hypothesized two-factor model (microjustice and macrojustice) fit the data reasonably well ($\chi^2 [n = 802, df = 19] = 123.53, p < .001$; CFI = .98; RMSEA = .08, CI = [.07, .10]; TLI = .97), and all items loaded significantly onto their intended factors. An alternative one-factor model fit the data poorly ($\chi^2 [n = 802, df = 20] = 1643.63, p < .001$; CFI = .64; RMSEA = .32, CI = [.31, .33]; TLI = .50). Consequently, we averaged the macrojustice items into one scale, which revealed high reliability (Cronbach’s $\alpha = .93$, coefficient H = .95), and the microjustice items into another scale, which also revealed high reliability (Cronbach’s $\alpha = .87$, coefficient H = .90).

We tested the content validity of these macrojustice and microjustice scales in Study S9 (see SM 2.9 for details). Specifically, we used the Hinkin and Tracey (1999) validation approach and the guidelines proposed by Colquitt, Sabey, Rodell, and Hill (2019). We interpreted these scores using Colquitt et al.’s (2019) evaluation criteria. We found that the microjustice and macrojustice scales had similar levels of definitional correspondence (“the degree to which a scale’s items correspond to the construct’s definition”) and definitional distinctiveness (“the degree to which a scale’s items correspond more to the construct’s definition than to the

definitions of other orbiting constructs” (Colquitt et al., 2019: 1243) as other commonly used organizational justice scales, namely distributive and procedural justice. The mean correspondence score for each scale was as follows: macrojustice ($M = .80, SD = .19$), microjustice ($M = .77, SD = .19$), distributive justice ($M = .87, SD = .15$), and procedural justice ($M = .83, SD = .12$). These levels are similar to the correspondence observed in other scales published in industrial and organizational psychology and organizational behavior (the scales examined in Colquitt et al.’s review of 112 scales published between 2010 and 2016 ranged from .60 to .96, $M = .87, SD = .06$).

The mean distinctiveness score for each scale was as follows: macrojustice ($M = -.01, SD = .07$), microjustice ($M = -.02, SD = .07$), distributive justice ($M = .03, SD = .06$), and procedural justice ($M = .01, SD = .06$). As Colquitt et al. (2019) explained, the range for the distinctiveness score is from -1 to 1, where a positive value reflects items receiving higher ratings on the intended construct than on orbiting ones, and a negative value reflects items receiving lower ratings on the intended construct than on the orbiting ones. The distinctiveness we observed also fell within the typical range. The scales examined in Colquitt et al. (2019) ranged from -.04 to .64 ($M = .27, SD = .14$). We note, however, that the distinctiveness scores may have been relatively low because we compared closely related orbiting constructs, which likely made it difficult for participants to distinguish between the scales’ items.

Additional measures. At the end of the study, participants answered questions that assessed how much their responses reflected socially desirable responding (i.e., the tendency for participants to respond in a way that will be viewed favorably by others; Paulhus, 1991). We used two abridged versions of common social desirability scales, the Marlowe–Crowne Social Desirability Scale (Crowne & Marlowe, 1960) and the Balanced Inventory of Desirable

Responding (Paulhus, 1991). We report analyses of these measures in SM 1.3 (we find that the social desirability scales improve prediction and do not account for the effect of decision type). At the end of the study, participants provided demographic information (age, gender, race or ethnicity, level of education, income, and political attitudes).

Results

Choice. Consistent with Hypothesis 1, participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (56.59%) than participants choosing between policies with equivalent consequences (45.58%; $\chi^2 [1, 802] = 9.64, p = .002$).

Likelihood. As with choice, participants reported a greater likelihood of choosing the higher-scoring, higher-income option when choosing between individuals ($M = 4.33, SD = 2.10$) than between equivalent policies ($M = 3.86, SD = 2.16; t[800] = 3.15, p = .002$).

Fair choice. Consistent with our theoretical account, participants also indicated that it was fair to choose the higher-scoring, higher-income option at a significantly higher rate when choosing between individuals (63.41%) than choosing between equivalent policies (38.67%; $\chi^2 [1, 802] = 48.70, p < .001$).

Microjustice and macrojustice. Participants indicated that their decision was (or would be, depending on the order of the measures) motivated by microjustice considerations significantly more when choosing between individuals ($M = 5.79, SD = 0.97$) than when choosing between equivalent policies ($M = 5.51, SD = 1.22; t[800] = 3.63, p < .001$). Conversely, they indicated that their decision was (or would be) motivated by macrojustice considerations marginally less when choosing between individuals ($M = 3.89, SD = 1.93$) than when choosing between equivalent policies ($M = 4.14, SD = 1.91; t[800] = 1.80, p = .072$).

Mediation analysis. We also tested whether the data were consistent with the proposed causal model in which the effect of making an individual versus policy decision on choosing the higher-scoring, higher-income option was mediated by one's macrojustice and microjustice concerns (Hypothesis 2; see Figure 1.1). Assuming this causal model, we conducted mediation analysis setting the choice of applicant or policy as the dependent variable (1 = higher-scoring, higher-income option, 0 = lower-scoring, lower-income option), the type of decision as the independent variable (1 = individual, 0 = policy), and the microjustice and macrojustice scales as simultaneous mediators. We ran the model using the PROCESS Macro for SPSS (Model 4, Hayes, 2018), which allowed us to estimate the indirect effect using a percentile bootstrap estimation approach with 10,000 samples.

We found that participants making the individual decision, as compared to the policy one, were significantly more motivated by microjustice concerns ($b = 0.28, p < .001$), which were significantly associated with a higher likelihood of choosing the higher-scoring, higher-income option ($b = 0.70, p < .001$; microjustice indirect effect = 0.20, 95% CI = [0.08, 0.33]). Furthermore, participants making the individual decision, compared to the policy decision, were marginally less motivated by macrojustice concerns ($b = -0.25, p = .072$), which were significantly associated with a lower likelihood of choosing the higher-scoring, higher-income option ($b = -1.15, p < .001$). However, due to the marginal effect of condition on macrojustice, the confidence intervals for the indirect effect did not omit zero (macrojustice indirect effect = 0.28, 95% CI = [-0.03, 0.60]). The indirect effects partially account for the relationship between making an individual versus a policy decision and decision-makers' choices. For completeness, we also depict the correlations between decision type, choice, macrojustice, and microjustice in Table 1.1.

Discussion

In Study 3, we tested Hypothesis 2, whether a relative reliance on microjustice versus macrojustice accounts for the relationship between making individual versus policy decisions and the likelihood of favoring higher-achieving, advantaged individuals over lower-achieving, disadvantaged individuals. We found that decision-makers' microjustice concerns, but not their macrojustice concerns, significantly mediated the relationship between making a policy decision versus an individual decision and their choice of applicant or policy. Overall, these results are consistent with Hypothesis 2. Indeed, in the individual condition, participants prioritized microjustice over macrojustice more than they did in the policy condition, although the absolute prioritization of macrojustice was only marginally different across conditions (interaction between judgment type [microjustice vs. macrojustice] and condition: $p = .002$). Furthermore, participants' prioritization of both microjustice and macrojustice predicted choice.⁶

We note, however, that our evidence of mediation in this study was also affected by whether the microjustice and macrojustice questions came before or after making a choice. Specifically, we found support for both microjustice and macrojustice as mediators when participants made their choice first and then considered how much it was motivated by microjustice and macrojustice. However, when participants considered how much their decision would be motivated by those concerns before making a choice, we did not find evidence of mediation through either microjustice or macrojustice (see SM 1.3 for details of these analyses). The naturally occurring differences between policy and individual decisions may have been

⁶ Notably, in other studies where we tested Hypothesis 2—albeit with earlier, unvalidated measures of microjustice and macrojustice—we also found evidence of mediation through both microjustice and macrojustice (see Study 4 and Study S7 in SM 2.7) and of an association between microjustice, macrojustice, and choice (see Study S8 in SM 2.8).

muted by participants considering the macrojustice and microjustice questions before making a choice. In other words, explicitly considering the hypothesized mechanism for the decisions prior to making a choice may have made the choices for each decision more alike. Indeed, the gap between policy and individual decisions itself, which is well-documented in our other experiments, was not significantly different when participants made their choice after the macrojustice and microjustice items. This attenuation of the gap may be particularly likely when participants consider macrojustice, which is typically less focal than microjustice (as evidenced by the mean levels of each) but more predictive of choice. We test this possibility directly—whether focusing participants on macrojustice concerns across conditions attenuates the policy–people gap—in Study 4.

We also found that the policy–people gap replicated overall even when policies and individual decisions had identical consequences. These results suggest that the mere notion that one is making a policy decision, as opposed to just a decision between two individuals, can be enough to create a discrepancy in choice and fairness between the two decision contexts. This result suggests that the policy–people gap cannot be explained by structural differences between the decisions, such as the number of people they affect and the identifiability of those people. Because the decisions had equivalent consequences, this gap also cannot be explained by previous accounts of discrepancies between the endorsement of broader principles and their implementation, such as decision-makers disagreeing with specific implementations of their principles (Dixon et al., 2007) or principles appearing to be less zero-sum than individual decisions (Ballinger & Crocker, 2020; Read & Loewenstein, 1995; Ritov & Zamir, 2014).

Study 4: An Intervention to Align Choices

In Study 4, we tested an intervention to align policy and individual choices (Hypothesis 3). The intervention had decision-makers learn about both macrojustice and microjustice standards of fairness and asked them to make decisions that prioritized the macrojustice standard. We predicted that asking decision-makers to focus on the macrojustice standard would affect individual decisions more than policy decisions because policy decisions are already motivated by macrojustice to a greater extent than individual decisions are. Because the intervention sought to manipulate the hypothesized mechanism underlying the gap between individual and policy decisions, it also allowed us to test Hypothesis 2 through moderation. Although the results of Study 3 implicate microjustice and macrojustice as potential mediators, they do not actually establish a causal relationship. Study 4 addressed this limitation, thus providing stronger evidence of the causal pathway between the type of decision, macrojustice and microjustice, and decision-makers' choices. We preregistered this study on AsPredicted.org (https://aspredicted.org/HIV_PJN).

Methods

Participants. We recruited online participants on Prolific Academic for a small monetary sum. We estimated the sample size we would require in order to be powered to detect a significant two-way interaction where the policy–people gap was attenuated by one-third. With effect sizes informed by our previous studies and pilots (i.e., $d = 0.9$ in the control condition vs. $d = 0.6$ in the intervention condition), we determined we would need 1,400 participants (where power = 0.80 and $\alpha = .05$). Since we required a large sample size, the only criteria for inclusion were that participants were based in the United States and had at least a college degree (so that they were familiar with college admissions). As indicated in our preregistration, we excluded

participants who were not able to see a sample table of information (as in Study 3), failed one of five attention and comprehension checks, or who answered the three questions about macrojustice and microjustice in under 6 seconds.⁷ Two checks were about the scenario and the information presented; these were identical to Study 3's checks. The remaining three checks sought to ensure that participants paid attention to the intervention and understood its content: "Were you asked to prioritize one of the views of fairness?" ("Yes, the micro view of fairness," "Yes, the macro view of fairness," or "No, I was simply asked to be conscious of them"), "Which of these best describes the micro view of fairness?" ("It is focused on the allocation of admission spots to individual applicants" or "It is focused on the overall allocation of admissions spots"), and "Which of these best describes the macro view of fairness?" (the same answer options were provided as in the preceding check). Our final sample comprised 1,346 participants ($M_{\text{age}} = 34.39$, $SD = 12.58$; 48.66% female).

Procedure and materials. We randomly assigned participants to a condition from a 2 (decision type: individual or policy) x 2 (intervention: control or macrojustice) between-subjects design. As in Studies 1 and 3, we asked all participants to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. In all conditions, we presented participants with the same table of information as in Study 1 and counterbalanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option was labeled "A" and appeared on the left.

⁷ For robustness, we also conducted our analyses including the 61 participants (4% of the sample) who answered the questions in under 6 seconds. All results were practically identical. For example, in the control condition, participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (54.48% with these participants vs. 54.69% without them) than participants choosing between policies (29.11% vs. 28.92%; $p < .001$ for both).

The control and macrojustice versions of the individual and policy conditions were identical to the individual and policy conditions tested in Study 1 (without the within-subjects portion), except they also had participants read about macrojustice and microjustice views of fairness and receive guidance on how to make their decisions. Before seeing any information about the individual applicants or policies, participants were told that the “Office of the University President has developed guidance for making admissions decisions.” All participants then read a description of the macrojustice and microjustice views of fairness. Control participants were told that “both of these standards are justifiable, but they are also imperfect,” so they “should simply be conscious of these two views as [they] make whatever admissions decision [they] deem appropriate.” Macrojustice participants were told that “at our university, we prioritize the macro view” of fairness and that they “should make admissions decisions that help accomplish a fair allocation of admission spots.” We depict the full text of the control and macrojustice conditions in Figure 1.3.

After reading the guidance about how to make admissions decisions and examining the table of information about the applicants or policies, participants selected one of the two applicants or policies (forced choice: A or B), indicated which applicant or policy was the fairer one to select (forced choice: A or B), and reported how likely they were to pick either option on a scale of 1 (*much more likely to choose A*) to 7 (*much more likely to choose B*). They also reported how much their decision was motivated by concerns related to the macrojustice and microjustice views of fairness, using an abbreviated scale. Specifically, on a scale from 1 (*not at all*) to 7 (*extremely*), participants reported how much their decision was motivated by two macrojustice concerns (“creating equality” and “increasing diversity”) and one microjustice concern (“rewarding applicants for their accomplishments”). We averaged the macrojustice

items into a composite ($r = .82$). Although these items were not validated (unlike those in Study 3), they tap into core aspects of microjustice and macrojustice and are similar to Study 3's items. Participants also provided demographic information (age, gender, race or ethnicity, level of education, income, and political attitudes).

Results

Choice. In the control condition, participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (54.69%) than participants choosing between policies (28.92%; $\chi^2 [1, 716] = 48.32, p < .001$). In the macrojustice condition, participants choosing between individuals also chose the higher-scoring, higher-income option at a significantly higher rate (41.18%) than participants choosing between policies (26.90%; $\chi^2 [1, 630] = 14.10, p < .001$), but this difference was smaller than in the control condition. We depict these results in Figure 1.4.

Consistent with our preregistration, to test whether the macrojustice intervention significantly attenuated the gap between the individual and policy decisions, we conducted an ordinary least squares (OLS) regression⁸ on the choice of applicant or policy (1 = higher-scoring, higher-income option; 0 = lower-scoring, lower-income option), using decision type (1 = individual, 0 = policy), intervention (1 = macrojustice, 0 = control), and the decision type x intervention interaction as independent variables. The coefficient on decision type was significant ($b = 0.26, p < .001$), indicating that participants in the control condition who made a choice about individuals were significantly more likely to select the higher-scoring, higher-income option than participants who made a choice about policies (54.69% vs. 28.92%). The

⁸ Per the recommendations of Gomila (2021), we conducted an OLS regression, rather than a nonlinear regression, to estimate treatment effects on binary outcomes. Nevertheless, we repeated this analysis with a logistic regression and found the same pattern of results.

coefficient for the intervention was not significant ($b = 20.02, p = .597$), indicating that the policy decisions in the control and macrojustice conditions (28.92% vs. 26.90%) were not significantly different. The absence of an effect of the intervention on the policy decisions suggests that policy decisions are naturally motivated by macrojustice concerns. In contrast, we found a significant interaction of decision type 3 intervention ($b = 20.12, p = .027$), indicating that the intervention shifted choices significantly more for decisions about individuals (54.69% vs. 41.18%) than policies (28.92% vs. 26.90%). This result also indicates that the macrojustice intervention brought individual decisions closer in line with policy decisions.

Likelihood. As with choice, in the control condition, participants reported a greater likelihood of choosing the higher-scoring, higher-income option when choosing between individuals ($M = 4.24, SD = 1.95$) than between policies ($M = 3.11, SD = 1.89; t[714] = 7.81, p < .001$). In the macrojustice condition, participants also reported a greater likelihood of choosing the higher-scoring, higher-income option when choosing between individuals ($M = 3.78, SD = 2.02$) than between policies ($M = 3.11, SD = 1.97; t[628] = 4.19, p < .001$), but this difference was smaller than in the control condition. We conducted the same OLS regression as with choice, and we found the same pattern of results: a significant decision type coefficient ($b = 1.13, p < .001$), a null intervention coefficient ($b = 20.00, p = .997$), and a significant decision type x intervention coefficient ($b = 20.46, p = .033$).

Fair choice. As with choice, in the control condition, participants choosing between individuals indicated that the higher-scoring, higher-income option was the fair choice at a significantly higher rate (59.90%) than participants choosing between policies (27.11%; $\chi^2 [1, 716] = 77.44, p < .001$). The control results were similar to those in the macrojustice condition (individual: 53.24% vs. policy: 28.28%; $\chi^2 [1, 630] = 40.09, p < .001$), but the difference between the control

and macrojustice conditions was not as pronounced as with choice. An OLS regression with the same format as the choice regression supports this pattern: a significant decision type coefficient ($b = 0.33, p < .001$), a null intervention coefficient ($b = 0.01, p = .759$), and a null decision type x intervention coefficient ($b = 20.08, p = .132$). The null interaction coefficient here suggests that the macrojustice intervention did not shift participants' judgments about which choice was the fair one to make.

Mediation analysis. We also tested whether the data were consistent with the proposed causal model in which the effect of making an individual versus policy decision on choosing the higher-scoring, higher-income option was mediated by one's macrojustice and microjustice concerns (Hypothesis 2; see Figure 1.1). Additionally, we tested whether having participants rely on the same standard of fairness would attenuate the difference between the individual versus policy decisions on choosing the higher-scoring, higher-income option. Assuming our proposed causal model, we conducted binary moderated mediation analysis, setting the choice of applicant or policy as the dependent variable (1 = higher-scoring, higher-income option, 0 = lower-scoring, lower-income option), the type of decision as the independent variable (1 = individual, 0 = policy), the macrojustice composite and the single microjustice item as simultaneous mediators, and the intervention as a moderator (1 = macro, 0 = control). We ran the model using the PROCESS Macro for SPSS (Model 8, Hayes, 2018), which allowed us to estimate the indirect effects using a percentile bootstrap estimation approach with 10,000 samples. Control participants making the individual decision, as compared to the policy one, were significantly more motivated by microjustice concerns ($b = 0.35, p = .002$), which were significantly associated with a higher likelihood of choosing the higher-scoring, higher-income option ($b = 0.84, p < .001$; control microjustice indirect effect = 0.30, 95% CI = [0.11, 0.50]). Unlike in

Study 3, control participants making the individual decision, as compared to the policy decision, were also significantly less motivated by macrojustice concerns ($b = -0.81, p < .001$), which were associated with a lower likelihood of choosing the higher-scoring, higher-income option ($b = 20.87, p < .001$; control macrojustice indirect effect = 0.70, 95% CI = [0.47, 0.99]).

For macrojustice participants, the indirect effect of macrojustice was smaller than for the control participants, yet it still omitted zero (macro macrojustice indirect effect = 0.37, 95% CI = [0.12, 0.63]; index of moderated mediation = 2.34, 95% CI = [-0.70, -0.01]). The indirect effect of microjustice was also smaller than for control participants, but it did not omit zero (macro microjustice indirect effect = 20.02, 95% CI = [-0.24, 0.20]; index of moderated mediation = 2.32, 95% CI = [-0.62, -0.04]). The indices of moderated mediation omitted zero, suggesting that the macro intervention attenuated the gap in choice between individual and policy decisions because it reduced—though it did not completely eliminate—the degree to which decision-makers’ standards of fairness shifted across the decision contexts. For completeness, in Table 1.2 we also depict the correlations between decision type, choice, macrojustice, and microjustice, both overall and splitting the data based on the presence of the intervention.

Discussion

As in the previous studies, participants who made decisions about individuals favored the higher-scoring, higher-income applicant at a higher rate than participants who made decisions about policies. In support of Hypothesis 2, we found that the choice difference between the individual and policy decisions in the control condition was mediated by the degree to which participants were motivated by microjustice and macrojustice. We also found that this difference was attenuated when participants were asked to make decisions that helped accomplish their institution’s priority of a macrojustice standard of fairness, providing causal evidence for our

theoretical model. Specifically, this led decision-makers choosing between individuals to select the lower-scoring, lower-income applicant at a higher rate, thus bringing choices about individuals closer in line with decisions about policies. Notably, the intervention did not shift policy decisions, which suggests that policy decisions were already naturally motivated by macrojustice concerns.

This study also provided support for Hypothesis 3 and, in doing so, showed the promise of a simple, face-valid intervention to align policy and individual choices. The intervention tested here simply had participants learn about macrojustice and microjustice and then asked them to prioritize the macrojustice standard of fairness. This was enough to reduce the policy–people gap by about half, so this intervention may be an attractive medium for organizations to improve the alignment between their selection policies and the individual decisions those policies intend to influence. That said, this intervention did not shift participants’ views about the fair choice to make, which is likely beyond the scope of a low-touch intervention like this one.

General Discussion

The present research documents and explains a gap between organizational selection policies and individual selection decisions. Decision-makers choose policies that do not favor the same applicants that they choose when deciding between specific applicants. Instead, decision-makers are more likely to choose selection policies that favor lower-achieving, disadvantaged applicants over higher-achieving, advantaged applicants than they are to select specific lower-achieving, disadvantaged applicants. We found this gap among 15 out of 16 preregistered studies, with samples of admissions officers responsible for making admissions decisions (Study 1), employees in technology occupations (Study 2), and lay people (Studies 3-4). We also found support for our theoretical account of the policy–people gap. We theorized that when people

make (a) individual selection decisions, compared to policy decisions, they rely more on microjustice relative to macrojustice and that (b) a greater relative reliance on microjustice than macrojustice is associated with favoring higher-achieving, advantaged applicants over lower-achieving, disadvantaged ones. Studies 3 and 4 support this account, both through mediation analysis and by manipulating decision-makers' reliance on macrojustice. Overall, our results strongly implicate a differential reliance on microjustice and macrojustice as explaining, at least in part, the policy–people gap.

We also developed an intervention that helped align choices about policies and individuals. In Study 4 we had decision-makers learn about both macrojustice and microjustice standards of fairness and then make decisions that helped accomplish their institution's priority of macrojustice. This simple, transparent intervention roughly halved the size of the policy–people gap, thus providing an easy way for organizations to try to align their selection policies and individual selection decisions. While the intervention did help align choices, it did not align decision-makers' underlying views about fairness, so it may not be effective at helping mitigate conflict around how one should make these decisions. An alternative intervention might be more effective, but decision-makers' views about the appropriate standard of fairness for each decision may also not be very flexible. However, it is encouraging that decision-makers are able to override these beliefs and choose consistently with the standards put forth by their organizations, especially in the case of organizations seeking to increase the representation of disadvantaged individuals in college and the workplace. This is particularly important because a series of fair individual decisions, in the microjustice sense, can lead to an aggregate result that is not seen as fair in the macrojustice sense.

Supplementary Findings

We document the robustness of the policy–people gap across multiple supplementary studies. Specifically, we find that the policy–people gap replicates across a wide range of applicant profiles (Studies S1 and S3); when we describe policies more naturalistically (Study S3) or as rules (Study S6); and with different forms of advantage and disadvantage including gender (Studies 2 and S5), high school graduation rate (Study S3), and household income (all other studies). The gap also persists in both between- and within-subjects designs (Study 1) and when we compare both single and repeated decisions (Study S2). See Table 1.2 for the key findings of each study we conducted and see the supplementary materials for further detail on all studies and additional analyses.

We also conducted an internal meta-analysis of all preregistered studies that tested Hypothesis 1 in order to identify a precise effect size estimate, which can be a useful benchmark for future researchers examining factors that systematically influence selection decisions. A random-effects model, which included a total of 8,653 participants, revealed a significant effect size for decision type on choice (odds ratio = 2.46, 95% CI = [2.02, 3.00], $z = 8.99$, $p < .001$, $Q(15) = 65.58$, $p < .001$). This result shows that, overall, the odds of favoring the lower-achieving, disadvantaged option were 2.46 times higher when choosing policies than when choosing individuals; these are equivalent to the odds of favoring the higher-scoring, advantaged option when choosing individuals than when choosing policies.⁹

⁹ The magnitude of the policy-people gap was similar across the manuscript studies (odds ratio = 2.27, 95% CI = [1.48, 3.46], $z = 3.79$, $p < .001$, $Q(3) = 14.79$, $p = .002$; $n = 2,058$) and the supplementary studies (odds ratio = 2.54, 95% CI = [2.04, 3.17], $z = 8.33$, $p < .001$, $Q(11) = 43.83$, $p < .001$; $n = 6,595$).

The effect size also reveals that the magnitude of the policy–people gap is relatively large. To contextualize its magnitude, we can compare it to research that has examined the selection of marginalized candidates. In eminent work on race and hiring decisions, Bertrand and Mullainathan (2004) found that the odds of an applicant with a White-sounding name receiving a call back was 1.55 times higher than of an applicant with a Black-sounding name. Another prominent example concerns the impact of blind auditions in remedying gender-biased hiring: the odds of a woman passing an orchestra audition’s preliminary round was 1.67 times higher when the audition was blind versus not blind (Goldin & Rouse, 2000). In the case of hiring for technology positions, the odds of hiring a woman were 1.05 times higher when researchers prompted people to increase the number of candidates on their hiring shortlist versus when people relied on their original shortlist (Lucas, Berry, Giurge, & Chugh, 2021). The overall size of the policy–people gap we identify is large in relation to these examples. Curiously, the most pared down version of a policy decision we tested (in Study 3), where selecting a policy had identical consequences as choosing an individual, created a gap of a similar magnitude as these examples (odds ratio = 1.56, 95% CI = [1.18, 2.06]). (See Figure 1.5, for each study’s odds.) In general, these comparisons stress the importance of bridging the policy–people gap because of the relatively large impact it has on people’s decisions about who to select.

Across our studies (specifically, Studies S2, S4–S7, and S10–S13), we also explored a range of alternative mechanisms for the policy–people gap and found limited support for them. For example, we did not find support that the policy–people gap could be explained as a conflict between what one wants to do and what one thinks one should do (Milkman et al., 2008), by differences in beliefs about the relationship between (dis)advantage and achievement across the decisions (Critcher & Dunning, 2013; O’Laughlin & Malle, 2002), or by differences in the

degree to which the decisions are perceived as zero-sum (Ballinger & Crocker, 2020). In contrast, we not only found support for the role of microjustice and macrojustice in explaining the gap (Studies 3-4), but also for our broader theorizing that policy decisions are construed more abstractly than individual decisions. To test for potential differences in construal level, we examined participants' textual explanations ($n = 3,725$) for why they selected the applicant and policy that they did (we have these data in Studies 2, S1, S3–S6, S10, S12, and S13). We used an online construal scoring tool (Humphreys, Isaac, & Wang, 2021) to measure the construal level reflected in participants' responses. We found that participants' explanations for their policy decisions reflected a more abstract construal than their explanations for individual decisions (see SM 4 for details).

Theoretical and Practical Contributions

Theoretically, this work identifies a novel gap between whom decision-makers select through their policies and through their specific selection decisions, and it provides a theoretical account for why this gap occurs. Decision-makers focus on what is fair to do across both policy and individual decisions, but they rely on different standards of fairness. This account provides a novel explanation for the known gaps between principles and their implementation, such as those concerning racial equality and diversity (e.g., Bell & Hartmann, 2007; Dixon et al., 2007; Smith & Mayorga-Gallo, 2017; Yogeeswaran & Dasgupta, 2014).

In contrast to existing work suggesting that decision-makers fail to think through the implementation of their principles (Bell & Hartmann, 2007; Smith & Mayorga-Gallo, 2017) or are influenced by temptation or cognitive blind spots when making isolated decisions (Bazerman & Tenbrunsel, 2011; Milkman et al., 2008), our account explains how policies and individual decisions can diverge, even when decision-makers do think carefully about their choices and

choose what they think is right. The relatively more abstract nature of policies (and principles, more broadly) shifts how much decision-makers rely on different standards of fairness when deciding what is right. Our theoretical account advances our understanding of ethical decision-making, fairness, and preference reversals in addition to contributing to the limited psychological research to date that has examined how decision-makers think about and create policies (e.g., Evers, Inbar, Blanken, & Oosterwijk, 2016; Nakashima, Daniels, & Laurin, 2017; Ritov & Zamir, 2014). Our work reveals how and why a common organizational tool—a selection policy—can lead to discrepant decisions relative to the subordinate decisions that a policy intends to influence.

In contrast to existing work suggesting that decision-makers fail to think through the implementation of their principles (Bell & Hartmann, 2007; Smith & Mayorga-Gallo, 2017) or are influenced by temptation or cognitive blind spots when making isolated decisions (Bazerman & Tenbrunsel, 2011; Milkman et al., 2008), our account explains how policies and individual decisions can diverge, even when decision-makers do think carefully about their choices and choose what they think is right. The relatively more abstract nature of policies (and principles, more broadly) shifts how much decision-makers rely on different standards of fairness when deciding what is right. Our theoretical account advances our understanding of ethical decision-making, fairness, and preference reversals in addition to contributing to the limited psychological research to date that has examined how decision-makers think about and create policies (e.g., Evers, Inbar, Blanken, & Oosterwijk, 2016; Nakashima, Daniels, & Laurin, 2017; Ritov & Zamir, 2014). Our work reveals how and why a common organizational tool—a selection policy—can lead to discrepant decisions relative to the subordinate decisions that a policy intends to influence.

This work also highlights the importance of studying macrojustice and microjustice within organizational contexts and facilitates their future application by developing valid scales to measure them. While these concepts have been used to examine affirmative action policies (Bobocel, Son Hing, Holmvall, & Zanna, 2002; Zdaniuk & Bobocel, 2011), to our knowledge, they have not been used to examine other organizational and public policy allocations (e.g., scarce medical resources, visas, and promotions). Indeed, it still appears true today, just as it was 30 years ago, that there is “a shortage of scholarship aimed at understanding how macro-level concerns (including societal-level inequalities) shape allocative processes within organizations and, conversely, how inequalities and injustices within organizations influence the broader society” (Baron & Cook, 1992: 196).

Though organizational scholars have long been focused on the construct of justice, those investigations have primarily focused on procedural and distributive justice (for reviews, see Colquitt et al., 2001; Colquitt et al., 2013). These constructs have greatly furthered our collective understanding of a range of organizational outcomes, including job satisfaction (McFarlin & Sweeney, 1992), job performance (Colquitt, LePine, Piccolo, Zapata, & Rich, 2012), and support for organizations and institutions (e.g., Sunshine & Tyler, 2003). However, these forms of justice are less useful for understanding how people conceptualize fairness at varying levels (e.g., for society versus individuals). Even decisions that are judged to be procedurally and distributively fair may lead to distributions that are not judged to be fair at a macro level. For example, even though decision-makers might agree that college admissions procedures are made free of individual bias and based on accurate information (procedurally fair), as well as in line with applicants’ performance and effort (distributively fair), they may nevertheless judge the resulting distribution of admitted applicants to be unfair, per their macrojustice concerns. To understand

situations such as this one, macrojustice and its counterpart, microjustice, should be considered. Moreover, the constructs of macrojustice and microjustice may be particularly helpful to study situations in which systemic inequality and structural differences between groups (e.g., in skill-building opportunities, encouragement, or inclusion) create differences in merit, thereby creating a tension between what is fair to individuals (per a microjustice concern for merit) and in society or organizations (per several possible macrojustice concerns, like equality or diversity). We facilitate future efforts to leverage these constructs by presenting validated scales of microjustice and macrojustice within selection contexts.

Practically, we uncover insights that can help organizations improve the correspondence between their selection policies—a ubiquitous and necessary organizational tool—and the subordinate decisions they intend to influence. Understanding why these decisions diverge in college admissions and workplace hiring can help create interventions to bridge the gap and thus better accomplish organizational and societal goals, such as increasing diversity and reducing inequality. Mitigating the gaps in these contexts is particularly important because they represent key moments of social and economic stratification. To that end, we tested a theory-based intervention that asked decision-makers to rely on their organization’s macrojustice standard of fairness. This straightforward intervention reduced the policy–people gap by about half and led to a greater number of selection decisions favoring disadvantaged individuals. We hope this intervention and its underlying theory is of use to practitioners and scholars alike.

Limitations and Future Directions

While this research found support for the proposed hypotheses, it did so with several limitations. One limitation is that we relied on vignette studies to document the policy–people gap. While this approach allowed us to control the differences between the policy and individual

decisions, it is low in external validity compared to, for example, examining policy and individual decisions in situ. However, additional recall studies with college admissions officers (Study S15) and management faculty (Study S14) revealed that the vignettes we used across studies accurately reflect real trade-offs that actual decision-makers face. Relatedly, we restricted our operationalization of policies to fit within this research approach. We held constant the information between policy and individual decisions in order to cleanly compare them. This afforded us careful experimental control at the expense of external validity. Nevertheless, we are encouraged by the results of Studies S2 and S3, where the policy–people gap replicates with richer, more naturalistic implementations of individuals and policies. Future work should examine the policy–people gap in more naturalistic settings with more externally valid stimuli.

A second limitation is that our studies did not directly test or measure differences in construal level, only differences in the reliance on microjustice and macrojustice. However, the exploratory analysis of participants’ explanations for their choices (see the Supplementary Findings section) is consistent with our theorizing. This is a promising preliminary finding, but more research is necessary to test whether construal level is indeed implicated, as we theorize. Investigating this question would not only test our theoretical account but also potentially shed light on a current debate about the effects of construal level on fairness and values (Alper, 2020).

Another limitation is that our studies rely on only a subset of possible macrojustice and microjustice principles. Therefore, while we propose that macrojustice and microjustice standards of fairness underlie the policy–people gap, it is possible that only egalitarian principles of macrojustice and meritocratic principles of microjustice underlie the gap, rather than macrojustice and microjustice as a whole. This ambiguity exists in our studies, as well as in other investigations of fairness standards that rely on specific operationalizations (e.g., distributive

justice is frequently operationalized as an equity principle; Colquitt, 2001), because meritocratic and (sometimes) egalitarian considerations are the relevant ones in most workplace contexts. Future research examining macrojustice and microjustice should consider contexts in which macrojustice and microjustice align with other principles (e.g., where microjustice aligns with need, as may be the case in aid and helping decisions). Studying microjustice and macrojustice across a broader range of contexts has the potential to further clarify their role in the policy–people gap and in other decision-making contexts.

One avenue for future research into the policy–people gap is to understand how personal characteristics influence it. Across our studies, we collected a number of individual-level demographic factors, namely, gender, age, race, income, education, and political attitudes. The policy–people gap was not clearly moderated by any of these factors except education. For example, for college admissions decisions that highlighted socioeconomic disadvantage and advantage, highly educated decision-makers were more likely to favor the higher-scoring, advantaged person when choosing between individuals, leading to a larger policy–people gap. We encourage future researchers to further examine how and why personal characteristics influence the policy–people gap.

Conclusion

This paper provides robust evidence for a policy–people gap, whereby decision-makers favor different individuals when making decisions about policies versus decisions about specific individuals. Specifically, decision-makers select policies that favor lower-achieving, disadvantaged individuals, but when selecting specific individuals, they favor higher-achieving, advantaged ones. This gap occurs because policy decisions shift the standards of fairness that motivate decision-makers' choices. However, the gap between decision-makers' policy and

individual choices can be attenuated by asking decision-makers to be consistent with their organization's standard of fairness. Overall, this paper helps us understand why individual selection decisions diverge from espoused organizational policies and suggests one way to close this gap, therefore helping to increase the representation of disadvantaged groups in organizations.

Chapter 2:
How Interpretations of Socioeconomic Advantage and Disadvantage Influence College
Admissions

Abstract

Why are college campuses not more socioeconomically diverse? Through 13 experiments and a simulation, we find that the way in which college applicants are evaluated upon learning about their socioeconomic circumstances skews admitted college classes towards the socioeconomically advantaged. Experimental participants evaluated a college applicant before and after learning that the applicant was either socioeconomically advantaged or disadvantaged. Fewer people adjusted their evaluations when the applicant was advantaged versus disadvantaged. This asymmetry in adjustment followed from people's beliefs about what socioeconomic advantage and disadvantage reveal about the effort applicants have to exert to achieve equivalent outcomes. In a simulation, we illustrate how this asymmetric adjustment leads to admitting a less socioeconomically representative class, thereby revealing one reason why college campuses may struggle to diversify. The supplementary materials linked at the end include details of the experiments.

Introduction

The socioeconomic makeup of the United States is skewed at selective college campuses. At “Ivy-Plus” colleges, there are more students from the top one-percent of the income distribution than from the entire bottom half (Chetty, Friedman, Saez, Turner, & Yagan, 2020). This skew exists despite the vast majority of colleges’ commitment to having a diverse student body (Starck, Sinclair, & Shelton, 2021).

Why are students from lower class backgrounds underrepresented in selective colleges and those from upper class backgrounds overrepresented? Social science research suggests several explanations for this inequality, including discrimination; disparities in access to human, social, and cultural capital; differences in accurate information about college; and rising tuition costs (Bastedo & Jaquette, 2011; Carnevale & Strohl, 2010; Chetty, Hendren, Kline, & Saez, 2014; Hoxby & Turner, 2015; Rivera, 2012; Stephens, Markus, & Phillips, 2014). We propose an additional explanation – that the way in which college applicants are evaluated in light of their socioeconomic circumstances skews admitted college classes towards the advantaged. There are two parts to this explanation.

The first part of this explanation is that people asymmetrically adjust their evaluations of socioeconomically advantaged versus disadvantaged applicants. We conducted 13 experiments in which people evaluated an applicant based only on their achievements and then again after learning that the applicant was either socioeconomically advantaged or disadvantaged. This two-stage process allowed us to examine how people adjusted their evaluations – or did not – in response to the socioeconomic information. Moreover, this process mirrors the way in which many admissions officers, as part of a widespread “holistic” approach to college admission, evaluate applicants in light of their circumstances (Bastedo et al., 2022; Bastedo, Bowman,

Glasener, & Kelly, 2018; Rigol, 2003; Sukumar & Metoyer, 2018; Sukumar, Metoyer, & He, 2018; College Board, 2016). On average, people adjusted their evaluation negatively based on advantage information and positively based on disadvantage information. However, they adjusted less for advantage than for disadvantage, creating an asymmetry.

We theorize that people asymmetrically adjust for dis/advantage because of their beliefs about how socioeconomic dis/advantage affects applicants and, thus, their corresponding inferences about applicants' effort. Most people adjust their evaluations of disadvantaged applicants positively because they believe that socioeconomic disadvantage makes success harder for applicants, thus requiring more effort. However, for socioeconomic advantage most people adjust negatively or do not adjust. Those who adjust negatively believe that advantage makes success easier, thus requiring less effort. Those who do not adjust believe that benefitting from advantage requires effort, which they already assumed of applicants. As a result of disadvantage having one modal direction of adjustment (positive), but advantage having two (negative and none), the average adjustment for advantage is smaller than the average adjustment for disadvantage.

The second part of this explanation is that asymmetric adjustment for dis/advantage can hinder admitting a socioeconomically representative college class. Because the primary factors in admissions evaluations are academic achievement (e.g., test scores) and academic achievement is positively correlated with socioeconomic circumstances, advantaged students are more likely to be admitted to college and disadvantaged students are less likely (Bastedo & Jaquette, 2011; Chetty, Friedman, Saez, Turner, & Yagan, 2020; Clinedinst, 2020; Reardon, 2016; Sackett, Kuncel, Arneson, & Waters, 2009). Adjusting evaluations based on applicants' circumstances can reduce this gap in admissions, thus creating a more representative college class (Bastedo et

al., 2022; Bastedo, Bowman, Glasener, & Kelly, 2018). However, we posit that the extent to which this gap is reduced depends on whether people adjust for both socioeconomic advantage and disadvantage. To test this, we simulated the outcome of college admissions cycles where the proportion of admissions officers adjusting their evaluations of advantaged and disadvantaged applicants varied. When fewer officers adjusted for advantage than for disadvantage (as we found in our experimental results), the admitted college class was less representative than when as many officers adjusted for advantage as for disadvantage. These results reveal how the socioeconomic composition of an admitted class is ultimately linked to admissions officers' beliefs about how socioeconomic dis/advantage affects applicants.

Asymmetric evaluations of socioeconomic advantage and disadvantage

Admissions Officer Survey

We ran a pilot survey in which we asked professional college admissions officers ($N = 35$) to imagine evaluating a college applicant with average qualifications. The officers then learned that the applicant was either socioeconomically advantaged or disadvantaged (all officers were asked about both, in random order), and reported whether the new information would change their impression of the applicant (1 *Not at all*, 4 *Moderately*, 7 *Extremely*). On average, the officers reported that they would somewhat adjust their evaluation if the applicant was advantaged ($M = 2.91$, $SD = 1.62$), but would adjust significantly more if the applicant was disadvantaged ($M = 4.14$, $SD = 1.46$; $t[34] = 4.24$, $p < .001$, $d = 0.72$). However, not all officers thought similarly about advantage and disadvantage information. While 26% of officers reported that their impression would change equally for advantage and disadvantage, 63% reported that they would adjust less for advantage than disadvantage. These results suggest that for most officers, advantage information is less relevant for updating their evaluations than disadvantage

information is, but that there is also a sizeable group for whom advantage and disadvantage are equally informative.

Experiments Examining Adjustment

To more granularly understand how and why socioeconomic dis/advantage information affects evaluations of an applicant, we experimentally manipulated the socioeconomic information that people learned about an applicant and examined how evaluations changed based on that information. In 12 experiments, we asked people to simulate the role of an admissions officer evaluating a college applicant seeking admissions to their competitive university.¹ We first presented participants with the applicant's academic achievement (e.g., SAT score, GPA) and asked them to evaluate the applicant. Participants decided whether to admit the applicant using a rating scale like those used in actual admissions processes, ranging from "Unqualified; definitely reject" to "Exceptional; definitely admit".² Participants also gave their impressions of the applicant's effort and talent.³ We then presented them with information indicating that the applicant was either socioeconomically advantaged or disadvantaged, and asked them again to evaluate the applicant and give their impressions. This two-stage process allowed us to examine the direction (positive or negative) and magnitude (none, a little, or a lot) of participants' adjustments upon learning about the applicant's circumstances.

¹ Detailed methods and results for all experiments are in the Supplementary Materials (SM).

² For example, a step in Princeton University's undergraduate admission process involves "summariz[ing] the applicant's achievements on the reader card, and not[ing] a suggestion action on the reader card: 'Admit,' 'Strong Interest,' 'Only if Room,' or 'Unlikely'" (Crosby, Iyer, Clayton, & Downing, 2003).

³ Though our theorizing focuses on effort, we measured both talent and effort because performance evaluations are typically attributed to both (Harrison, Kravitz, Mayer, Leslie, & Lev-Arey, 2006; College Board, 2016).

The socioeconomic information about the applicant described an advantageous or disadvantageous high school and neighborhood (Figure 2.1). The content and presentation of the information was based on the College Board’s Landscape™ tool, which is increasingly used in U.S. college admissions as part of holistic evaluation practices (Bastedo et al., 2022; Bastedo, Bowman, Glasener, & Kelly, 2018). We designed the disadvantage and advantage information to be symmetric in their intensity. For example, the median household income and per pupil funding for disadvantage were at the 25th and 17th percentiles, whereas for advantage were at the 75th and 83rd percentiles. Moreover, we pretested our stimuli to ensure that: (a) people did not differentially expect that the applicant was dis/advantaged before seeing the information and (b) that the disadvantage information was not seen as more disadvantaging than the advantage information was seen as advantaging.

Admissions evaluations. We synthesized the results of all 12 experiments through a pre-registered internal meta-analysis. On average, people adjusted their admissions evaluations positively for disadvantage ($d = 0.74$, 95% $CI = [0.41, 1.08]$) and negatively for advantage ($d = -0.42$, 95% $CI = [-0.52, -0.32]$); $N = 1541$, $K = 3$; Figure 2.2, Panel A). Consistent with our hypothesis, their adjustments were asymmetric: people adjusted positively for disadvantage information more than they adjusted negatively for advantage information. To further understand why people adjusted in this way, it is necessary to decompose their average adjustments into *how much* they adjusted and *the direction* in which they adjusted. These two components – magnitude and direction – can create an asymmetry in adjustment, but each suggests different reasons for adjusting.

There was no difference in the magnitudes of adjustment for advantage versus disadvantage. Figure 2.2 (Panel B) shows the effect size for the difference in the magnitudes.

Conditional on adjusting, the confidence interval for the difference in the magnitudes does not omit zero ($d = 0.15$, 95% $CI = [-0.11, 0.41]$, $z = 1.13$, $p = .257$). In other words, among people who adjusted, they did not adjust more or less for disadvantage versus advantage.

In contrast, the directions in which people adjusted created the asymmetry (Figure 2.2, Panel C). Most people adjusted positively for disadvantage (59.97%), but people were split between negative adjustment (42.52%) and no adjustment (38.23%) for advantage. In other words, people were more likely to adjust than not adjust for disadvantage than for advantage (*Odds Ratio [OR]* = 1.82, 95% $CI = [1.46, 2.26]$, $z = 5.34$, $p < .001$) and more likely to adjust positively for disadvantage than negatively for advantage ($OR = 2.03$, 95% $CI = [1.66, 3.49]$, $z = 6.82$, $p < .001$). These results reveal that the adjustment asymmetry stemmed from people adjusting for advantage in one of two ways, versus primarily one way for disadvantage. To understand why people adjusted their evaluations in this way, we also examined people's impressions of the applicant's effort and talent.

Effort and talent impressions. People adjusted their impressions of the applicant's effort as they did the admissions evaluations.⁴ They adjusted their effort impressions positively for disadvantage ($d = 0.87$, 95% $CI = [0.65, 1.09]$), negatively for advantage ($d = -0.54$, 95% $CI = [-0.71, -0.38]$), and more for disadvantage than advantage. Again, the adjustment magnitudes were not asymmetric ($d = -0.03$, 95% $CI = [-0.24, 0.17]$, $z = 0.34$, $p = .734$), but the directions were. Most people adjusted their effort impressions positively based on disadvantage information

⁴ We also examined impressions of warmth (see SM1). We expected minimal adjustment because the socioeconomic information is not pertinent to these impressions. As expected, most people did not adjust their impressions in both the disadvantage (50.38%) and advantage (57.30%) conditions.

(66.47%), but were mixed based on advantage information: some adjusted negatively (46.79%) and others did not adjust (42.13%).

Adjustments for talent impressions were less asymmetric than for effort. On average, people adjusted positively for disadvantage ($d = 0.35$, 95% $CI = [0.22, 0.48]$) and negatively for advantage ($d = -0.23$, 95% $CI = [-0.36, -0.09]$), but the absolute difference between these effects and their confidence intervals was smaller than for effort. Moreover, unlike with effort, people adjusted positively for disadvantage (44.63%) about as often as they did not adjust (42.75%). For advantage, people mostly adjusted negatively (30.75%) or did not adjust (54.76%). As with effort, the adjustment magnitudes were not asymmetric ($d = -0.01$, 95% $CI = [-0.19, 0.17]$, $z = 0.13$, $p = .893$).

Effort was more central than talent in people's evaluation of the applicant. We asked people to explain their admissions evaluation, and then we coded their explanations for mentions of effort-related (e.g., "effort," "hard work," and "work ethic") and talent-related words (e.g., "competence," "intelligence," and "smartness"). People mentioned effort-related words (43.2%) significantly more than talent-related words (10.3%, $z = 12.36$, $p < .001$). Because effort-related words were more common than talent-related ones, and effort adjustments were also more common than talent adjustments, we theorized that people's beliefs about the applicant's effort in particular might explain the overall asymmetry.

Theory about socioeconomic dis/advantage and effort inferences

Decades of psychological research suggest that people conceptualize the effects of advantage as *tailwinds* and of disadvantage as *headwinds* (Davidai & Gilovich, 2016; Jones, 1990; Phillips & Lowery, 2020; Reeder, Hesson-McInnis, Krohse, & Scialabba, 2001).

Tailwinds are helpful because they make a task easier, like running with the wind at one's back,

whereas headwinds are detrimental because they make a task more difficult, like running against the wind. This way of thinking about dis/advantage leads to straightforward inferences about effort. For example, kicking a soccer ball with the wind at one's back is seen as requiring less effort than kicking the ball that same distance with the wind blowing in one's face (Reeder, Hesson-McInnis, Krohse, & Scialabba, 2001). The effort adjustments we observe, whereby people adjust negatively for a college applicant's advantage and positively for their disadvantage, suggest that people do also reason about socioeconomic dis/advantage as *winds*. However, we contend that this is not the only way people reason about socioeconomic circumstances, particularly advantage.

We posit that people also conceptualize the effects of socioeconomic advantage as a *well-tuned bicycle*. A well-tuned bicycle is helpful because it allows one to move forward more efficiently than on a poorly-tuned bicycle. However, one still needs to continuously turn the pedals to move forward or. This way of thinking about advantage suggests that effort is still required to reap the benefit of the advantage. Hence, a college applicant who attended a private, well-resourced school be seen would be seen as having benefitted from attending the school (e.g., their test score is higher because of it), but the applicant's effort is not discounted because they had to put in effort to benefit.

In the case of socioeconomic disadvantage, the distinction between a headwind and a poorly-tuned bicycle is superfluous. A person is seen as having to exert more effort in order to achieve the same outcome as someone without that disadvantage, whether the disadvantage is understood as a headwind—a force that makes a task more difficult—or as a poorly-tuned bicycle—an inefficient vehicle for translating effort into forward progress.

We theorize that the direction in which people adjust their effort impressions follows from their beliefs about how socioeconomic dis/advantage affects applicants. Hence, we expect most people to positively adjust their evaluation when learning about an applicant's *disadvantage*. However, we expect people to either negatively adjust (following a “winds” model) or not adjust (following a “well-tuned bicycle model”) their evaluations when learning about an applicant's *advantage*. In other words, there is one dominant direction of adjustment in response to disadvantage, but two in response to advantage.

Ample research finds that people's beliefs about social class are associated with their political attitudes (Bullock, Williams, & Limbert, 2003; Davidai & Gilovich, 2018; Davidai & Ongis, 2019; Jost, Federico, & Napier, 2009; Kraus & Tan, 2015; Kteily, Sheehy-Skeffington, & Ho, 2017). Hence, if people beliefs underlie how they adjust (as we theorize), then we would also expect their political attitudes to be associated with their adjustment. Figure 2.3 shows how people's political attitudes (from the 12 experiments in the meta-analysis) relate to how they adjusted their impressions of the applicant's effort. Most participants across the political spectrum adjusted positively for disadvantage. However, for advantage, more liberal participants adjusted negatively, whereas moderate and more conservative participants did not adjust. While these results are consistent with our theorizing, they do not specifically point to people's beliefs about socioeconomic circumstance as the source of why people adjust asymmetrically. To more directly test our theorizing, we conducted an additional experiment.

Experiment Testing Theory

We experimentally manipulated people's beliefs about what socioeconomic advantage suggests about applicants' effort ($N = 732$, preregistered).⁵ We hypothesized that shifting people's beliefs about advantage in line with a 'winds' model – i.e., that advantage propels people forward, requiring less effort to reach a specific outcome -- would lead more people to negatively adjust their effort impressions. In contrast, shifting people's beliefs in line with a 'bicycle' model – i.e., that advantage helps people move efficiently, but still requires effort to reach a specific outcome – would lead fewer people to adjust. We randomly assigned participants to one of four conditions: Disadvantage, Advantage-Control, Advantage-Wind, and Advantage-Bicycle. The Disadvantage and Advantage-Control conditions were identical to those in our previous experiments, in which people made an initial rating based only on the applicant's academic achievements and then a subsequent rating after learning about the applicant's circumstances through the Landscape tool.

Wind vs. bicycle mental models. In the Wind and Bicycle conditions, we sought to influence participants' beliefs by explaining (a) how socioeconomic advantage is either like having the wind at one's back or like riding a bicycle and (b) what this implies about how to interpret the applicant's socioeconomic information. We then measured participants' beliefs by asking them how much they agreed (1=*Strongly Disagree*, 7=*Strongly Agree*) with three statements ($\alpha = .71$) about the relationship between effort and socioeconomic advantage (e.g., "Socioeconomically advantaged applicants need to work hard in order achieve high scores and

⁵ We ran a similar preregistered experiment with different stimuli. However, our stimuli in this experiment failed to effectively manipulate beliefs about whether socioeconomic advantage necessitates effort, precluding us from testing our hypothesis.

accomplishments” and “Effort is always needed to reap the benefits of socioeconomic advantage”).

The manipulation shifted participants’ beliefs: Compared to Control participants ($M = 4.72$, $SD = 1.22$), Wind participants believed that socioeconomic advantage necessitates effort less ($M = 4.32$, $SD = 1.38$, $t[374] = 3.01$, $p = .002$), whereas Bicycle participants believed that it necessitates effort more ($M = 5.03$, $SD = 1.09$, $t[380] = 2.57$, $p = .011$).

Participants’ beliefs about advantage were correlated with their overall effort adjustment ($r = .31$). However, only Wind participants significantly adjusted their effort impressions compared to Control participants ($M = -0.95$, $SD = 1.19$ vs. $M = -0.19$, $SD = 1.38$; $t[374] = 5.67$, $p < .001$); the adjustments of Control and Bicycle participants did not significantly differ ($M = -0.39$, $SD = 1.24$; $t[380] = 1.49$, $p = .137$). This pattern is reflected in the directions in which participants adjusted (Figure 2.4, Panel A). Wind participants adjusted differently than Control ($\chi^2[2, 376] = 27.89$, $p < .001$) and Bicycle participants ($\chi^2[2, 396] = 21.50$, $p < .001$); Control and Bicycle participants did not adjust differently ($\chi^2[2, 382] = 1.19$, $p = .552$).

The proportion of participants who adjusted negatively supports our theorizing. In the Wind condition, 65.64% of participants adjusted negatively compared to less than half of Control (44.20%; $\chi^2[1, 376] = 17.46$, $p < .001$) and Bicycle participants (45.77%; $\chi^2[1, 396] = 15.83$, $p < .001$). In other words, when participants more strongly believed that benefitting from socioeconomic advantage necessitates effort (as in the Wind condition) more of them adjusted negatively.

Contrary to our expectations, there was no differences in rates of non-adjustment. We expected more participants to not adjust when they less strongly believed that advantage does not necessitate effort (as in the Bicycle condition). However, only slightly more participants did not

adjust in the Bicycle condition (32.34%) than in the Control (29.28%, $\chi^2[1, 382] = 0.42, p = .519$) and Wind conditions (26.67%, $\chi^2[1, 396] = 1.53, p = .216$). These results suggest that participants' beliefs about advantage are just one of several factors that influence whether participants adjust or not.

Admissions evaluations. Participants adjusted their admissions evaluations in line with their effort adjustments. Indeed, adjustments to the admissions evaluations were correlated with the overall effort adjustments ($r = .60$), as well as with beliefs about advantage ($r = .25$). Only Wind participants adjusted their admissions evaluations significantly differently from the Control participants ($M = -0.55, SD = 1.23$ vs. $M = -0.29, SD = 1.30$; $t[374] = 2.00, p = .046$); the adjustments of Bicycle and Control participants did not significantly differ ($M = -0.23, SD = 1.18$; $t[380] = 0.42, p = .674$). The directions of adjustment bore the same pattern as with effort: Marginally more participants adjusted negatively in the Winds condition compared to the Control ($\chi^2[1, 376] = 3.00, p = .083$) and Bicycle conditions ($\chi^2[1, 396] = 5.25, p = .022$). In terms of not adjusting, 42.29% of participants did not adjust in the Bicycle condition, compared to 36.46% in the Control condition ($\chi^2[1, 382] = 1.35, p = .245$) and 35.90% in the Wind condition ($\chi^2[1, 376] = 0.01, p = .909$). These results suggest that effort impressions and admissions evaluations are closely related and that shifting advantage beliefs influences not only effort impressions, but also evaluations about who should attend college.

Shifting Wind participants' admissions evaluations eliminated the asymmetry between advantage and disadvantage. As our meta-analysis reveals, people adjust less for advantage than for disadvantage. However, in this experiment Advantage-Control and Disadvantage participants adjusted to similar extents ($M = -0.29, SD = 1.22$ vs. $M = 0.23, SD = 1.45$). Nevertheless, we are able to ascertain the effect of shifting participants' adjustment for advantage. On average, Wind

participants adjusted negatively ($M = -0.59$, $SD = 1.23$) twice as much as Advantage-Control participants did ($t[374] = 2.00$, $p = .046$). These results reveal that shifting participants' beliefs about advantage in line with a Winds model can mitigate or even reverse the asymmetry.

Across our studies, we found that people asymmetrically evaluated socioeconomically advantaged versus disadvantaged applicants. The way people adjusted their evaluations was related to their impressions of the applicants' effort, which in turn followed from their beliefs about how socioeconomic dis/advantage affected applicants. Ahead, we focus on the consequences of such asymmetric adjustment. We simulate how asymmetric adjustment for advantage versus disadvantage can hinder admitting a socioeconomically representative college class.

Asymmetric evaluations based on dis/advantage preclude diverse college classes

We simulated the outcome of a college admissions cycle if officers adjusted their impressions of disadvantaged and advantaged applicants symmetrically versus asymmetrically. To do so, we created a dataset of hypothetical applicants representing the entire income distribution in the U.S (Donovan, Dalaker, Labonte, & Romero, 2021). We assigned those applicants an SAT score percentile sampled from a distribution of scores for their income group (College Board, 2016; Zhang, 2020). We simulated which applicants would be admitted under different evaluation criteria (Figure 2.5, Panel B). First, we simulated no adjustment based on applicants' income. Here, applicants were admitted based solely on their SAT scores, an important component of their academic achievement.

Next, we simulated symmetric adjustment (Figure 2.5, Panel C). This simulation represents a situation where all officers hold a winds model of dis/advantage, resulting in symmetric adjustment. The adjustments consisted of a two- to eight-point adjustment to an

applicant's SAT score percentile. If the applicant's income was below the \$80,000 threshold (i.e., the midpoint of the income categories and the 55th percentile), the adjustment was positive, and if the income was above the threshold the adjustment was negative. The adjustments increased from two points to four, six, and eight points for the income categories that were further away from \$80,000; this progressive adjustment reflects our assumption that people would adjust their impressions more the more dis/advantaged the applicant was. Hence, for example, applicants in the \$60,000-\$80,000 income category would receive +2 percentile points, whereas applicants in the \$100,000-\$140,000 would receive -4 percentile points.

Finally, we simulated asymmetric adjustment (Figure 2.5, Panel D), such that only 54% of advantaged applicants received an adjustment. This percentage comes from the ratio of participants in our meta-analysis who negatively adjusted their evaluations versus did not adjust, thus reflecting implications of our findings for real-world college admissions.

The no-adjustment simulation created the most unequal admitted class because it maintained the existing income inequality in SAT scores – for example, only 27% of admitted applicants had incomes below \$80,000 despite representing 55% of all applicants.⁶ Adjusting symmetrically led to an admitted class with a similar distribution of incomes as the pool of applicants (59% admitted with incomes below \$80,000). In contrast, adjusting asymmetrically—as we find people do in our experiments—led to a less representative admitted class (41% admitted with incomes below \$80,000). These results, along with the experimental studies, show that the way in which people evaluate applicants in light of their socioeconomic circumstances can lead to admitting a student body that under-represents disadvantaged applicants and over-

⁶ A similar “SAT only” simulation (Blanchard, 2015) using nationally representative datasets from the U.S. Department of Education yielded similar proportions (e.g., only 32% of applicants admitted to highly selective institutions had incomes below \$75,000).

represents advantaged applicants, even when people both have access to and try to integrate socioeconomic status information into their evaluations.

Implications

This work furthers our understanding of how people think about socioeconomic dis/advantage when evaluating applicants. Our theory and findings provide a more complete picture of how people conceptualize the effects of socioeconomic advantage – not only as having the wind at one’s back, but also as riding a well-tuned bicycle – and how this impacts the evaluation of applicants. These insights are timely as colleges increasingly use socioeconomic information in the evaluation of candidates (Bastedo et al., 2022), and possibly more so if race-conscious practices are deemed unconstitutional.

These findings further suggest that affirmative action policies and procedures, such as holistic admissions, may not necessarily lead to their intended outcomes. While ample work has considered who supports these kinds of initiatives and why (Crosby, Iyer, & Clayton, 2003; Crosby, Iyer, & Sincharoen, 2006; Harrison, Kravitz, Mayer, Leslie, & Lev-Arey, 2006; Ho, & Unzueta, 2015; Lowery, Unzueta, Knowles, & Goff, 2006), the present work shows how the effectiveness of these policies depends on who is implementing them. Specifically, the socioeconomic composition of an admitted class is linked to whether admissions officers believe that socioeconomic advantage reveals information about applicants’ effort or not. Although the accuracy of such beliefs is impossible to verify, officers nevertheless admit applicants based on them. Thus, if colleges aim to diversify their campuses, as many do for moral or instrumental reasons (Starck, Sinclair, & Shelton, 2021), then they may better meet that goal if their officers assume that socioeconomic advantage and disadvantage are equally informative about applicants.

Chapter 3:
Luck that Builds Merit

Abstract

A common way that luck affects people's outcomes is by giving them the opportunity to develop merit. We know that people generally find it unfair to allocate rewards based on luck, but fair to allocate them based on merit. However, we do not know whether people find it fair or unfair to allocate rewards based on luck that builds merit. This work seeks to understand how people judge fairness and how they allocate rewards in situations when luck builds merit. We find a gap between what people regard as fair and what they do. Rewarding people who were lucky to build merit is regarded to be less fair than rewarding people who simply have merit, seemingly with no influence of luck, and fairer than rewarding people who were simply lucky, but have no merit. Nevertheless, people reward those who were lucky to build merit as much as those who simply have merit. We also shed light on three judgments that influence how people judge the fairness of situations in which luck builds merit and how they reward them: whether the benefits conferred by luck are internalized in the form of skill, whether the benefits are lasting, and whether reaping those benefits required effort. These findings help understand why people find present inequalities to be unfair, but may nevertheless behave in ways that reinforce them.

“Through dumb luck, [my business partner] Charlie and I were born in the United States, and we are forever grateful for the staggering advantages this accident of birth has given us.”

—Warren Buffett, 2014 Letter to Berkshire Shareholders

Introduction

Luck plays a role in many outcomes people hope for. The financial situation of the country and family one is born into are strong predictors of one’s own income in adulthood (Bowles, Gintis, Osborne Groves, 2005; Chetty, Kline, Hendren, & Saez, 2014; Milanovic, 2015). Having facial features that suggest competence make one more likely to be elected for public office (Ballew & Todorov, 2007; Todorov, Mandisodza, Goren, & Hall, 2005). Having an earlier birth month increases the odds of becoming a professional hockey player (Addona & Yates, 2010). Having a more pronounceable name and earlier surname initials make one more likely to succeed in certain fields, like economics and law (Einav & Yariv, 2006; Laham, Koval, & Alter, 2012). The list goes on.¹ Is it fair that luck affects these outcomes?

An extensive body of research finds that people regard differences in outcomes due to luck to be less fair than those based on merit—whether defined as performance, production, ability, or effort (Alesina, Glaeser, & Sacerdote, 2001; Almås, Cappelen, Sørensen, Tungodden, 2010; Cappelen, Eichele, Hugdahl, Specht, Sørensen, & Tungodden, 2014; Cappelen, Konow, Sørensen, & Tungodden, 2013; Durante & Putterman, 2014; Krawczyk, 2010). These fairness perceptions are important – they predict whether people want to maintain or mitigate luck-based inequalities (Alesina & Angeletos, 2005; Bénabou & Tirole, 2006; Starmans, Sheskin, & Bloom, 2017). When these inequalities are perceived to be unfair, people seek to mitigate them by, for

¹ See, for example, Robert Frank’s 2016 book, *Success and Luck: Good Fortune and the Myth of Meritocracy*, or Scott Barry Kaufman’s 2018 *Scientific American* article, “The Role of Luck in Life Success Is Far Greater Than We Realized.”

example, allocating resources more equally (e.g., Tinghög, Andersson, & Västfjäll, 2017) or supporting redistributive measures, like progressive taxation (e.g., Durante & Putterman, 2014). However, this body of work has operationalized luck in ways that may not represent the world we live in or the kinds of judgments that people make about luck (e.g., luck that multiplies the product of one's labor, Almås, Cappelen, Sørensen, Tungodden, 2010; outcomes that directly follow from luck, without any participation, Cappelen, Konow, Sørensen, & Tungodden, 2013).

We posit that luck can often lead to the development of merit, specifically skill. Consider being born to a wealthy family. This lucky situation gives a person access to better and more schooling, artistic tutors, and sports coaches (Beamish, 1990; Borowiecki, 2019; Chetty, Friedman, Saez, Turner, & Yagan, 2020). The luck of being born to this family does not automatically translate into selection into college, musical groups, or sports teams. Rather, this luck gives the opportunity to develop the skills that will help one be selected for these. Although the circumstances of one's birth (in terms of one's family, neighborhood, country of residence, etc.) are arguably the most common, consequential, and clearest form of luck, the idea that luck builds merit is not specific to it. Luck can refer to any circumstance that is external to a person and outside their control. Hence, luck can build merit in other ways. For example, one can win the lottery to attend a selective charter school, thus developing better academic skills,² or have a serendipitous encounter that provides expert advice for improving one's skill.³

Luck that builds merit creates a tension. It pits the unfairness of luck with the fairness of merit. Imagine, for example, having to decide whether to admit a college applicant whose family

² <https://www.nydailynews.com/new-york/bronx/stakes-high-lottery-bronx-charter-school-98-families-home-unhappy-article-1.110796>

³ <https://www.indiatimes.com/sports/did-you-know-sachin-tendulkar-once-wanted-to-be-a-fast-bowler-yes-you-read-that-right-256660.html>

was able to pay for them to attend a top-performing high school, thus giving the applicant the opportunity to become well prepared for college. On the one hand, it seems unfair to reward this applicant for being lucky. Other students did not have this luck, and so it seems unfair to make a decision that rewards the applicant for a factor that was outside of their control. On the other hand, the applicant is in fact well-prepared for college—that is, they have merit—and they probably had to work hard to develop college readiness. It seems fair to reward them on the basis of that merit. This is how the same choice – to admit the applicant – can seem both unfair and fair. What do people actually think is fair to do in this situation and what do people actually do? We seek to answer these questions.

We hypothesize that rewarding people who were lucky to build merit is regarded to be (a) less fair than rewarding people who simply have merit, seemingly with no influence of luck, and (b) fairer than rewarding people who were simply lucky, but have no merit. Moreover, we expect that people reward those who were lucky to build merit as much as those who simply have merit (even though they think it is less fair to do so).

We also seek to understand *why* luck that builds merit may be perceived to be less fair than just merit, but more than just luck. To do so, we focus on three factors that distinguish situations in which luck builds merit from situations where people are simply lucky. The first of these factors is that luck becomes internalized. In situations in which luck builds merit, the benefit conferred by the luck becomes internalized in the form of skill. The second factor is that the benefit of the luck is lasting. The benefit conferred by the luck continues to be beneficial in the future. The third factor is that the luck requires effort to reap its benefit. In order to develop merit, the luck requires a greater exertion of effort. We independently manipulate these factors to examine their influence on people's judgments of fairness and who they choose to reward.

Experiment

We conducted an experiment that put people in the hypothetical position of having to evaluate players with different kinds of luck. This is akin to selection contexts like college admissions and workplace hiring, where applicants with differential luck compete for selection. This experiment is a precursor to one we intend to run where participants' decisions will actually lead to real consequences for others.

In this study we also sought to examine why situations in which luck builds merit might be deemed fairer than situations in which luck simply leads to better performance. Developing merit from luck might be deemed to be fairer because (1) the beneficiary of the luck is actually more skilled, (2) the benefit of the luck is lasting, and (3) effort is required to translate luck into merit.

Methods

We recruited 151 participants on Prolific Academic (Age: $M = 37.6$, $SD = 11.9$; Gender: 42.4% Female, 56.3% Male, 1.3% Non-Binary). Participants in this study evaluated hypothetical "Players". They first read that we had conducted a study where we had participants ("Players") complete Wordle puzzles and that we now needed help deciding which few Players to invite to a follow-up, high-paying Wordle tournament. They learned what the Players had done in the previous study: "Players attempted to solve 3 [Wordle](#) puzzles... Before attempting the Wordle puzzles, Players were randomly assigned to a 'situation.' We left it entirely up to chance which situation each Players would be assigned to. The situations affected Players in some way. You will learn more detail about each situation as you evaluate Players." We told participants that they would evaluate 10 Players and we showed them the type of information they would see about each Player. Participants then sequentially evaluated the 10 Players. For each Player,

participants saw information about their performance, effort, situation and predicted performance in the tournament. Table 3.1 shows the exact information we presented about each Player.

All participants evaluated 10 Players corresponding to a 2 (Effort: Low, High) x 2 (Internalization: Not Internal, Internal) x 2 (Lastingness: Limited, Lasting), plus 2 (Just Merit Low Effort, Just Merit High Effort) within-subjects design. Additionally, participants were randomly assigned to one of three between-subjects conditions (Player Performance: Low, Medium, High), which determined the Players' average number of guesses (5, 4, 3).⁴

We showed the Players in random order and labeled each with a randomly generated ID number (e.g., Player ID 542).

The within-subjects factors – Effort, Internalization, and Lastingness – determined the information that participants saw about each Players. The Effort factor determined how much effort – low or high – the Player exerted in solving the puzzles. We conveyed this information by indicating how much time the Player spent on the puzzles and how much effort they put in, based on their own self-reports.

The Internalization factor determined whether the Player's situation had provided them with an advantage that became internalized or not, that is, whether it had actually developed skill. We conveyed this information by describing two different types of situations. To convey advantage that became internalized, we described the Player as receiving training (e.g., learning the most common vowels and consonants in 5-letter words) to help them solve the puzzles in fewer guesses. To convey advantage that was not internalized, we described the Player as simply receiving some of the letters for each puzzle or, in other words, part of the answer.

⁴ Because this between-subjects factor does not interact with fairness or choice, for ease of exposition we do not consider it in our analyses.

The Lastingness factor determined whether the benefit provided by the Player's situation would continue to be helpful. To convey lasting benefits, we explained that the training would continue to be helpful because the tournament would draw on the same set of words or that the Player would continue to receive free letters in the tournament. In both of those cases we mentioned that Player would likely perform similarly to how they had performed so far. To convey limited benefits, we explained that the tournament would use different words than the training was based on or that the Player would no longer receive free letters, such that the Player would likely perform worse than how they had performed so far.

The Just Merit Players differed from the other Players in that we did not describe internalization or lastingness. For these Players, instead of providing information about training or free letters, we indicated that these Players "simply played puzzles" and that during the tournament they "will simply attempt more puzzles..." so they "will probably perform similarly to how they performed here." These Players' effort was conveyed in the same way as for the other Players.

Participants answered seven questions about each Player. The questions were presented on the same page and always in the order shown below:

- Effort: "How much effort would you say the Player put into solving the puzzles?"
(1=*No effort at all* to 5=*An extreme amount of effort*)
- Hard Work: "How hard did this Player have to work to solve the puzzles in [#] guesses?" (1=*Not hard at all* to 5=*Extremely hard*)
- Skill: "How skilled do you think this Player is at Wordle puzzles?" (1=*Not skilled at all* to 5=*Extremely skilled*)

- Internalization: “Did the Player's situation actually improve their skill at Wordle?”
(1=*Did not improve at all* to 5=*Extremely improved*)
- Prediction: “If this Player plays in the tournament, how many guesses will it take them to solve the puzzles (on average)?” (*1 guess, ..., 6 guesses, Will not be able to solve the puzzles*)
- Choice: “If you were paid a bonus based on how well this Player performed in the tournament, would you invite this Player to the tournament?” (1=*Definitely NOT*, 3=*Maybe*, 5=*Definitely YES*)
- Fairness: “Would it be unfair or fair to invite this Player to the tournament?”
(1=*Definitely NOT*, 3=*Maybe*, 5=*Definitely YES*)

After rating all Players, participants answered questions examining their understanding of the scenario. To assess whether participants actually perceived the training and free letters situations as being due to luck, we asked them whether being assigned to a particular situation was in or out of Players’ control (*Out of their control, Maybe in their control, In their control*). As intended, 86% of participants reported that it was out of Players’ control. This result is important because it establishes that most participants recognized the Internalization situations to be based on luck. To examine whether participants understood Players’ scores, we asked them whether it is better to have a higher or lower number of guesses (*Higher, Not Sure, Lower*); 91% correctly reported lower.

We also asked participants two exploratory questions to assess whether the effort information about the Player suggested that they were less talented. We asked whether learning that the Player spent more effort solving the puzzle made them think that the Player was less naturally talented. Participants were mixed: 49% said *Sometimes*, 28% said *No*, 22% said *Yes*,

and 2% said *Other*. We also asked an open-ended question about what the time information suggested about the Player. Participants' provided a range of interpretations reflecting similar inferences as the effort question (e.g., "they thought about their guess so they are probably more logical," "the more time spent made me think they weren't as good at solving the puzzles"). While these results likely introduce some noise to our manipulations, we do not expect them to limit our ability to test our hypothesis. Future experiments with this paradigm should try to minimize inferences about the applicants' talent based on their effort.

Results

We report the means and standard deviations for all Player measures in Table 3.2 and the correlations in Table 3.3. Ahead, we first describe the results for the choice and fairness measures, followed by the other evaluations of the Players.

Fairness. We depict participants' fairness judgments in Figure 3.1. Participants rated the Just Merit Players as fairest to invite. The difference between the high-effort and low-effort was significant (difference between means = .32, paired samples t-test: $t[150] = 3.34, p = .001$). To assess the differences between the other Players, we conducted a factorial repeated measures ANOVA with factors for the within-subjects variables: Effort, Internalization, and Lastingness. We observed significant main effects of Effort ($F[1, 150]=21.36, p < .001$), Internalization ($F[1, 150]=89.31, p < .001$), Lastingness ($F[1, 150]=5.69, p = .017$), and a significant interaction of Internalization and Lastingness ($F[1, 150]=13.78, p < .001$). These results reveal that effort is a relevant consideration for fairness – across the board, participants deemed it fairer to invite Players who exerted more effort to achieve their score (difference between means = .26, paired samples t-test: $t[150] = 3.63, p < .001$). In the case of Internalization, participants regarded it fairer to invite Players who internalized the benefit of their luck versus Players who did not

internalize it (difference between means = .53, paired samples t-test: $t[150] = 7.12, p < .001$). In the case of Lastingness, participants regarded it marginally fairer to invite Players whose benefit would last versus those whose benefit would not last (difference between means = .13, paired samples t-test: $t[150] = 1.90, p = .059$). However, the differences between Internalization and Lastingness were qualified by an interaction, such that the Internalization difference was larger when the benefits were lasting versus not lasting. There was a bigger difference between internalized and non-internalized luck between Players whose luck provided a lasting benefit (difference between means = .74, paired samples t-test: $t[150] = 4.26, p < .001$) than between Players whose luck did not grant a lasting benefit (difference between means = .32, paired samples t-test: $t[150] = 7.78, p < .001$).

Two comparisons are of key interest to our hypotheses: whether the Lasting Training, High-Effort Player was deemed as fair to invite as the Just Merit, High-Effort Player and as fair to invite as the Limited Letters, Low-Effort Player. The Lasting Training, High-Effort Player represents the situation in which luck develops skill that lasts and that takes effort to develop, which is arguably the most common situation in which luck builds merit. The Limited Letters, Low-Effort Player represents luck that lacks the key elements that distinguish luck that builds merit, as well as the kind of luck that has been represented in previous research (e.g., Almås, Cappelen, Sørensen, Tungodden, 2010; Cappelen, Konow, Sørensen, & Tungodden, 2013; Tinghög, Andersson, & Västfjäll, 2017). As hypothesized, inviting the Lasting Training, High-Effort Player was regarded as less fair than inviting the Just Merit, High-Effort Player (difference between means = .63, paired samples t-test: $t[150] = 5.81, p < .001$), but more fair than inviting the Limited Letters, Low-Effort Player (difference between means = .63, paired samples t-test:

$t[150] = 4.86, p < .001$). These results suggest that luck that builds merit is both regarded to be fairer than simply being lucky, but nevertheless discounted for the presence of luck.

Choice. We depict participants' choices in Figure 3.2. As noted before, the phrasing of the choice question intended to incentivize participants to choose the Players who they thought would perform best. Indeed, this choice measure was strongly correlated with participants' perceptions of the Players' skill ($r = .65$) and expectations of who would perform best in the tournament (i.e., solve the puzzles in the *fewest* guesses; $r = -.43$). Participants gave their highest invitation rating to the Just Merit, High-Effort Player ($M = 3.50, SD = 0.94$), as well as to the Lasting Training, High-Effort Player ($M = 3.50, SD = 1.06$; difference between means = .01, paired samples t-test: $t[150] = 0.09, p = .939$). Despite thinking that the Lasting Training, High-Effort Player was less fair to invite, people were just as likely to select them.

To assess the differences between the other Players, we again conducted a factorial repeated measures ANOVA with factors for the within-subjects variables: Effort, Internalization, and Lastingness. We observed significant main effects of Effort ($F[1, 150]=24.21, p < .001$), Internalization ($F[1, 150]=56.72, p < .001$), and Lastingness ($F[1, 150]=210.02, p < .001$), and no significant interactions. For the choice to invite Players, there was qualitatively-speaking a larger difference between Lastingness Players (difference between means = .80, paired samples t-test: $t[150] = 10.74, p < .001$), than between Effort Players (difference between means = .28, paired samples t-test: $t[150] = 3.49, p < .001$) and Internalization Players (difference between means = .42, paired samples t-test: $t[150] = 8.20, p < .001$). These results suggest that the decisions about whether to reward lucky individuals—when one is incentivized to choose those who will perform best—are particularly sensitive to the lastingness of luck's benefits, whereas

decisions about who is fair to reward are particularly sensitive to the internalization of luck's benefits.

Effort. We depict participants' effort judgments in Figure 3.3. Given the strong correlation between the effort and hard work measures ($r = .74$), we only focus on the effort measure. The patterns of results are not qualitatively different based on the measures. We conducted a repeated factorial repeated measures ANOVA with factors for the within-subjects variables: Effort, Internalization, and Lastingness. We found significant main effects of effort ($F[1, 150] = 1931.51, p < .001$) and Internalization ($F[1, 150] = 24.88, p < .001$), as well as a marginally significant interaction of Effort and Internalization ($F[1, 150] = 3.54, p < .001$). As intended by our manipulation of Players' effort, high-effort Players were seen as having exerted more effort than low-effort Players (difference between means = 1.54, paired samples t-test: $t[150] = 26.43, p < .001$). Additionally, Players who internalized the benefit of their luck were seen as having exerted more effort than those who did not internalize it (difference between means = 0.18, paired samples t-test: $t[150] = 4.67, p < .001$).

Skill. We depict participants' effort judgments in Figure 3.4. The pattern of results was the same as with choice. We conducted a repeated factorial repeated measures ANOVA with factors for the within-subjects variables: Effort, Internalization, and Lastingness. We observed significant main effects of Effort ($F[1, 150] = 12.36, p < .001$), Internalization ($F[1, 150] = 141.76, p < .001$), and Lastingness ($F[1, 150] = 9.35, p = .002$), and no significant interactions.

Internal. We depict participants' internalization judgments in Figure 3.5. The pattern of results was the same as with fairness. We conducted a repeated factorial repeated measures ANOVA with factors for the within-subjects variables: Effort, Internalization, and Lastingness. We observed significant main effects of Effort ($F[1, 150] = 16.47, p < .001$), Internalization ($F[1,$

150]=18.61, $p < .001$), and Lastingness ($F[1, 150]=55.42, p < .001$), and a significant interaction of Internalization and Lastingness ($F[1, 150]=4.31, p = .038$).

Tournament Prediction. We depict participants' tournament prediction in Figure 3.6. We conducted a repeated factorial repeated measures ANOVA with factors for the within-subjects variables: Effort, Internalization, and Lastingness. We observed significant main effects Internalization ($F[1, 150]=45.46, p < .001$) and Lastingness ($F[1, 150]=332.94, p < .001$), and a significant interaction of Internalization and Lastingness ($F[1, 150]=8.94, p = .003$). Participants expected that, among Players with lasting benefits, those who internalized their luck would perform slightly better than who did not (difference between means = .17, paired samples t-test: $t[150] = 3.33, p = .003$). However, among Players with limited benefits, they expected those who internalized their luck to perform better by a larger degree than those who did not (difference between means = .43, paired samples t-test: $t[150] = 6.40, p < .001$).

Discussion

We sought to examine luck that builds merit because of the tension it creates between the unfairness of rewarding luck and the fairness of rewarding merit. People regarded rewarding individuals who were lucky to build merit to be less fair than rewarding individuals who simply had merit. However, people were just as likely to reward both individuals. This separation between what people find to be fair and what they do can help to understand, for example, why wealth continues to be associated with more favorable outcomes in a range of contexts, such as college admission, employment, art, and sports, despite people regarding it as unfair.

People also regarded rewarding luck that builds merit to be fairer than simply rewarding luck. We found that this was the case largely because luck that builds merit internalizes the benefits conferred by luck. This finding helps us to better understand how people reason about

luck and when they may find it be more or less fair. Moreover, this finding brings nuance to the monolithic version of luck that has been extensively researched and to more recent work examining allocations when luck and merit are present in unknown quantities (Cappelen, Mollerstrom, Reme, & Tungodden, 2022; Cappelen, Moene, Skjelbred, & Tungodden, 2023). Specifically, our findings reveal that people care not just about whether luck leads to differential outcomes, but also about the way in which luck brings forth those outcomes.

Appendix A: Tables

Table 1.1: Correlations between Independent and Dependent Variables (Studies 3-4), Chapter 1

| Panel A. Study 3 | | | | | |
|-------------------------|--------|----------|----------|---|--|
| Variable | 1 | 2 | 3 | 4 | |
| 1. Decision Type | - | | | | |
| 2. Choice | 0.11** | - | | | |
| 3. Microjustice | 0.13** | 0.33*** | - | | |
| 4. Macrojustice | -0.06† | -0.67*** | -0.21*** | - | |

| Panel B. Study 4 | | | | | | |
|-------------------------------|----------|----------|---------|----------|----------|---|
| Variable | 1 | 1a | 1b | 4 | 5 | 6 |
| 1. Decision Type | - | | | | | |
| 1a. when Macro = Control | - | - | | | | |
| 1b. when Macro = Intervention | - | - | - | | | |
| 4. Choice | 0.21*** | 0.26*** | 0.15*** | - | | |
| 5. Microjustice | 0.05* | 0.12** | -0.01 | -.49*** | - | |
| 6. Macrojustice | -0.17*** | -0.22*** | -0.11** | -0.60*** | -0.31*** | - |

Note. $N = 802$ (Study 3) and $N = 1346$ (Study 4). Decision Type was coded "1" for the Individual condition and "0" for the Policy condition. Choice was coded "1" when selecting the higher-scoring, advantaged option and "0" when selecting the lower-scoring, disadvantaged option. Rows 1a and 1b of Panel B depict correlations between Decision Type and the other variables for each of the Macro conditions (Control and Intervention).

- † $p < .10$
- * $p < .05$
- ** $p < .01$
- *** $p < .001$

Table 1.2: Overview of Main and Supplementary Studies, Chapter 1

| Study | <i>n</i> | Main Finding / Purpose |
|------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | 190 | Admissions officers favored higher-scoring, advantaged applicants over lower-scoring, disadvantaged applicants more often when selecting individual applicants than selecting policies. |
| 2 ^a | 360 | Tech employees favored higher-scoring, male applicants over lower-scoring, female applicants more often when selecting individuals to interview than when selecting interview policies. |
| 3 | 802 | The policy-people gap established in Studies 1 and 2 replicated and was driven by reduced microjustice concerns when making policy decisions. |
| 4 | 1,346 | The policy-people gap was attenuated when decision makers were asked to prioritize the macrojustice standard of fairness. |
| S1 | 2,016 | The policy-people gap replicated with 24 out of 25 different combinations of applicant profiles. |
| S2 | 1,009 | The policy-people gap replicated when participants made multiple individual decisions versus a policy decision followed by multiple decision to fine-tune the policy. |
| S3 | 603 | The policy-people gap replicated describing policies broadly or specifically, as well as when the dis/advantage information was household income or high school graduation rate. |
| S4 | 807 | The feature of policies that most affected choice was the framing of a decision as one about policies; less influential was the number of individuals affected and their identifiability. |
| S5 ^a | 859 | The feature of policies that most affected choice was the framing of a decision as one about policies; less influential was the number of individuals affected. |
| S6 | 593 | The policy-people gap replicated when policies were described as "rules." |
| S7 | 283 | The policy-people gap replicated when policies had identical consequences as individual decisions, and was driven by greater macrojustice concerns when making policy decisions. This study relied on preliminary macrojustice and microjustice scales. |
| S8 | 240 | An initial (failed) development of macrojustice and microjustice scales. Making a policy versus individual decision did not influence these scales. |
| S9 | 149 | We developed macrojustice and microjustice scales that correspond to those constructs and are moderately distinct from other related constructs. (i.e., procedural and distributive justice). |
| S10 | 595 | An intervention that highlighted the adversity that applicants have faced attenuated the policy-people gap. |
| S11 | 595 | Instructing participants to rely on specific values when choosing individuals and to rely on objective information when choosing policies narrowed the policy-people gap. |
| S12 | 298 | We explored several alternative explanations of the policy-people gap. |
| S13 | 287 | We explored several alternative explanations of the policy-people gap. |
| S14 ^a | 108 | Management faculty report facing tradeoffs between job candidates' achievements and circumstances. Their focus on a macrojustice standard of fairness predicts their selection of a disadvantaged candidate over a candidate with slightly better objective metrics of achievement. |
| S15 | 35 | Admissions officers report facing tradeoffs between college applicants' achievements and circumstances. |

Note. All studies except S14 and S15 were preregistered. ^a The context for Studies 2, S5, and S14 was workplace hiring. The context for all other studies was college admissions.

Table 3.1: Overview of Studies, Chapter 1

| Experimental Condition | Information about Players Read by Participants |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Effort | <i>Average Time Spent on Puzzle + Player Effort Solving Puzzles</i> |
| High | 3 minutes 4/5 (high effort) |
| Low | 1 minute 2/5 (low effort) |
| Internal | <i>Situation</i> |
| Internal | "Expert Training." This Player received access to a training to teach them how to solve the puzzles in fewer guesses. The training taught them the most common vowels and consonants in 5-letter words and where those letters tend to come up. |
| Not Internal | "Free Letters." This Player was given some of the letters for each puzzle. In other words, they received part of the answer for each puzzle. |
| Lasting | <i>Tournament Prediction</i> |
| Lasting | (if internal) The training will continue to be helpful in the tournament (the tournament will use the same set of words that the training is based on), so the Player will probably perform similarly to how they performed here (if not internal) This Player WILL continue to receive more free letters in the tournament, so the Player will probably perform similarly to how they performed here |
| Limited | (if internal) The training will NOT be helpful for the words in the tournament (the tournament will use a different set of words than the training is based on), so the Player will probably perform worse than how they performed here (if not internal) This Player will NOT receive any more free letters in the tournament, so the Player will probably perform worse than how they performed here |

Note. Participants evaluated Players based on each combination of the following factors: 2 (Effort: High, Low) x 2 (Internal: Internal, Not Internal) x 2 (Lasting: Lasting, Limited). We randomized the order of the Players that participants saw.

Table 3.2: Means and Standard Deviations of Dependent Variables, Chapter 3

| Player | | | Choice | | Fairness | | Skill | | Internalization | | Prediction | | Effort | | Hard Work | |
|------------|--------------|--------|----------|-----------|----------|-----------|----------|-----------|-----------------|-----------|------------|-----------|----------|-----------|-----------|-----------|
| Lasting | Internal | Effort | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> |
| Lasting | Internal | High | 3.50 | 1.06 | 3.38 | 1.19 | 2.96 | 0.95 | 3.22 | 0.89 | 3.89 | 0.91 | 3.81 | 0.60 | 3.45 | 0.76 |
| | | Low | 3.11 | 1.23 | 3.14 | 1.20 | 2.79 | 0.98 | 3.02 | 1.02 | 3.91 | 1.06 | 2.21 | 0.70 | 2.13 | 0.77 |
| | Not Internal | High | 3.03 | 1.07 | 2.62 | 1.26 | 2.36 | 0.93 | 2.91 | 1.23 | 4.10 | 1.02 | 3.58 | 0.74 | 3.12 | 0.94 |
| | | Low | 2.75 | 1.18 | 2.41 | 1.13 | 2.26 | 0.85 | 2.67 | 1.28 | 4.03 | 1.07 | 2.06 | 0.58 | 2.00 | 0.85 |
| Limited | Internal | High | 2.63 | 1.09 | 3.36 | 1.09 | 2.79 | 0.93 | 2.76 | 1.11 | 4.60 | 1.17 | 3.80 | 0.61 | 3.42 | 0.74 |
| | | Low | 2.38 | 1.04 | 3.00 | 1.06 | 2.62 | 0.96 | 2.49 | 1.15 | 4.52 | 1.31 | 2.17 | 0.67 | 2.23 | 0.85 |
| | Not Internal | High | 2.19 | 1.04 | 2.97 | 1.17 | 2.30 | 0.89 | 2.58 | 1.27 | 4.99 | 1.36 | 3.54 | 0.80 | 3.14 | 0.95 |
| | | Low | 1.98 | 0.97 | 2.75 | 1.10 | 2.13 | 0.79 | 2.44 | 1.31 | 4.98 | 1.36 | 2.11 | 0.64 | 2.05 | 0.83 |
| Just Merit | | High | 3.50 | 0.94 | 4.01 | 0.89 | 3.17 | 0.81 | 1.92 | 1.04 | 3.99 | 0.91 | 3.81 | 0.55 | 3.47 | 0.78 |
| | | Low | 3.15 | 1.24 | 3.70 | 1.11 | 2.99 | 1.09 | 1.64 | 0.89 | 3.97 | 1.02 | 2.22 | 0.69 | 2.15 | 0.82 |

Note. $n = 151$. Participants evaluated 10 Players, corresponding to the combination of 2 (Lasting: Lasting, Limited) x 2 (Internal: Internal, Not) x 2 (Effort: High, Low) factors and two additional “Just Merit” Players with High and Low effort.

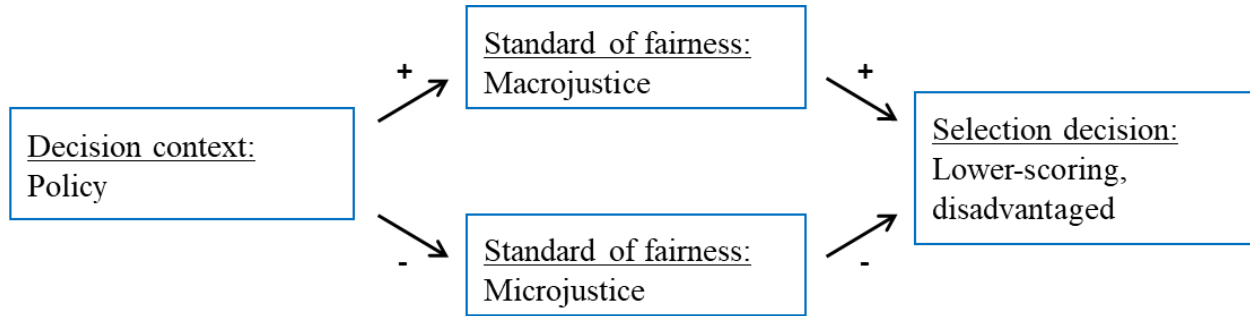
Table 3.3: Correlations between Dependent Variables, Chapter 3

| Variable | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|--------------------|----------|-------------------|----------|----------|--------------------|----------|----------|
| 1. Choice | - | | | | | | |
| 2. Fairness | 0.44 | - | | | | | |
| 3. Skill | 0.65 | 0.5 | - | | | | |
| 4. Internalization | 0.25 | 0.02 [^] | 0.27 | - | | | |
| 5. Prediction | -0.43 | -0.14 | -0.33 | -0.09 | - | | |
| 6. Effort | 0.21 | 0.23 | 0.22 | 0.15 | -0.03 [^] | - | |
| 7. Hard Work | 0.14 | 0.23 | 0.18 | 0.12 | -0.03 [^] | 0.74 | - |

Note. All factors significantly correlated at the $p < .001$ level unless indicated otherwise by [^], which denotes $p > .10$

Appendix B: Figures

Figure 1.1 Theoretical Model for the Link between Policy Decisions and Individual Selection Decisions, Chapter 1



Note. A decision about policies, compared to a decision about individuals, increases decision-makers' reliance on a macrojustice standard of fairness and decreases their reliance on a microjustice standard of fairness. Decision-makers' reliance on macrojustice and microjustice standards of fairness are positively and negatively associated, respectively, with selecting lower-scoring, disadvantaged applicants over higher-scoring, advantaged applicants.

Figure 1.2 Stimuli (Study 1), Chapter 1

Panel A. Individual Decision Table

| | Applicant A | Applicant B |
|---------------------------------------------------------|----------------------------|-----------------------------|
| SAT score | 1200 | 1350 |
| SAT percentile | 81st percentile | 94th percentile |
| GPA | 3.8 | 3.8 |
| Extracurriculars (#) | 3 | 3 |
| Letters of Recommendation | 2 strong recommendations | 2 strong recommendations |
| Estimated Household Income (percentile, approx.) | \$60,000 (50th percentile) | \$120,000 (80th percentile) |

Panel B. Policy Decision Table

| | Applicant type A | Applicant type B |
|---------------------------------------------------------|----------------------------|-----------------------------|
| SAT score | 1200 | 1350 |
| SAT percentile | 81st percentile | 94th percentile |
| GPA | 3.8 | 3.8 |
| Extracurriculars (#) | 3 | 3 |
| Letters of Recommendation | 2 strong recommendations | 2 strong recommendations |
| Estimated Household Income (percentile, approx.) | \$60,000 (50th percentile) | \$120,000 (80th percentile) |

Note. Participants used these tables of information to make their individual and policy decisions. Participants made both decisions, one after another (the first decision and corresponding table was on the first page of the survey and the second decision and table was on the second page). We randomized the order of the decisions across participants.

Figure 1.3 Stimuli (Study 4), Chapter 1

Panel A. Text Read by All Participants



Guidance on undergraduate admissions decisions

The Office of the University President has developed the following guidance for making fair undergraduate admissions decisions. There are limited spots available at our university and an increasing number of qualified applicants competing for these spots. We know that it is difficult to decide which applicants to admit and that there are multiple perspectives of how to make fair admissions decisions.

People can have different views on what is fair. When it comes to admitting college applicants, we consider that there are two main views of what it means to admit applicants fairly. One perspective—the micro view—is focused on what is fair to the applicants themselves, whereas the other perspective—the macro view—is focused on what is fair for society.

The micro view is focused on the allocation of admissions spots to individual applicants. Under this view, admissions decisions are fair if they are based on the merits of the applicants being evaluated, such as their accomplishments and qualifications. Fairness is defined as the degree to which the admissions process involves recognizing applicants' merits, and rewarding them in kind. This means that admissions officers should make admissions decisions that correspond to individual applicants' level of merit.

The macro view is focused on the overall allocation of admissions spots. Under this view, admissions decisions are fair if they lead to an ideal distribution of admitted students. Fairness is defined as the degree to which admissions decisions create an ideal distribution of college admission offers, for example, by creating equality across different groups, by accurately representing society as a whole, or by diversifying the current college landscape. This means that admissions officers should make admissions decisions that help accomplish an ideal allocation of admission spots in society.

Panel B. Additional Text Read Only by Control Participants

Guidance on undergraduate admissions decisions

We recognize that both of these standards are justifiable, but they are also imperfect.

This means that, as an admissions officer, you should simply be conscious of these two views as you make whatever admissions decisions you deem appropriate.

Panel C. Additional Text Read Only by Macrojustice Participants

Guidance on undergraduate admissions decisions

We recognize that both of these standards are justifiable, but they are also imperfect.

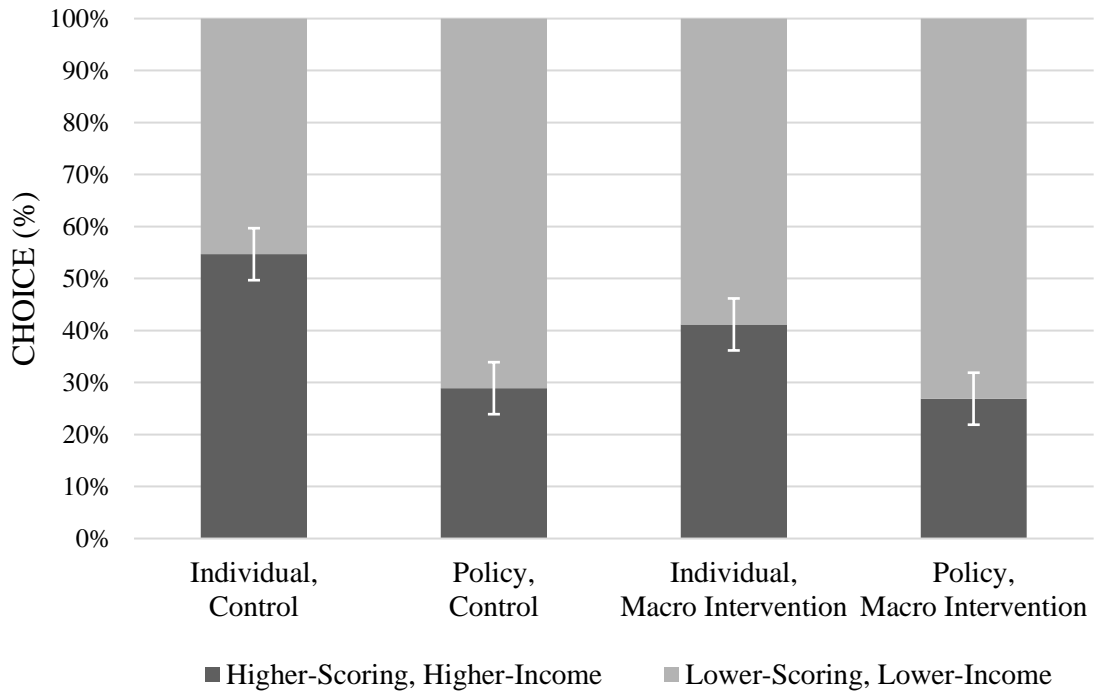
Nonetheless, at our university we prioritize the macro view.

We believe that barriers continue to exist in society that prevent college admissions from reflecting the ideal of equal opportunity. Therefore, we seek to create a class of admitted students that is fair by taking into account what the admitted class should look like.

This means that, as an admissions officer, you should make admissions decisions that help accomplish a fair allocation of admission spots.

Note. Participants in the control and macrojustice conditions read the introductory text (Panel A) before reading their condition-specific guidance (either Panel B or Panel C).


Figure 1.4 Choice by Decision Type and Intervention Condition, Chapter 1



Note. Error bars represent 95% confidence intervals (Study 4).

Figure 1.5 Forest Plot of Internal Meta-Analysis on Choice (Studies 1-4, S1-S8, S10-S13),

Chapter 1



Note. Odds ratios indicate the odds of favoring the lower-achieving, disadvantaged option when choosing policies versus when choosing individuals. Squares show effect-size estimates (odds ratios). Error bars show 95% confidence intervals (CIs). The diamond represents the meta-analytic effect size estimate and 95% CI. Weights are from a random-effects model.

Figure 2.1 Experimental Stimuli, Chapter 2

(A) Advantage

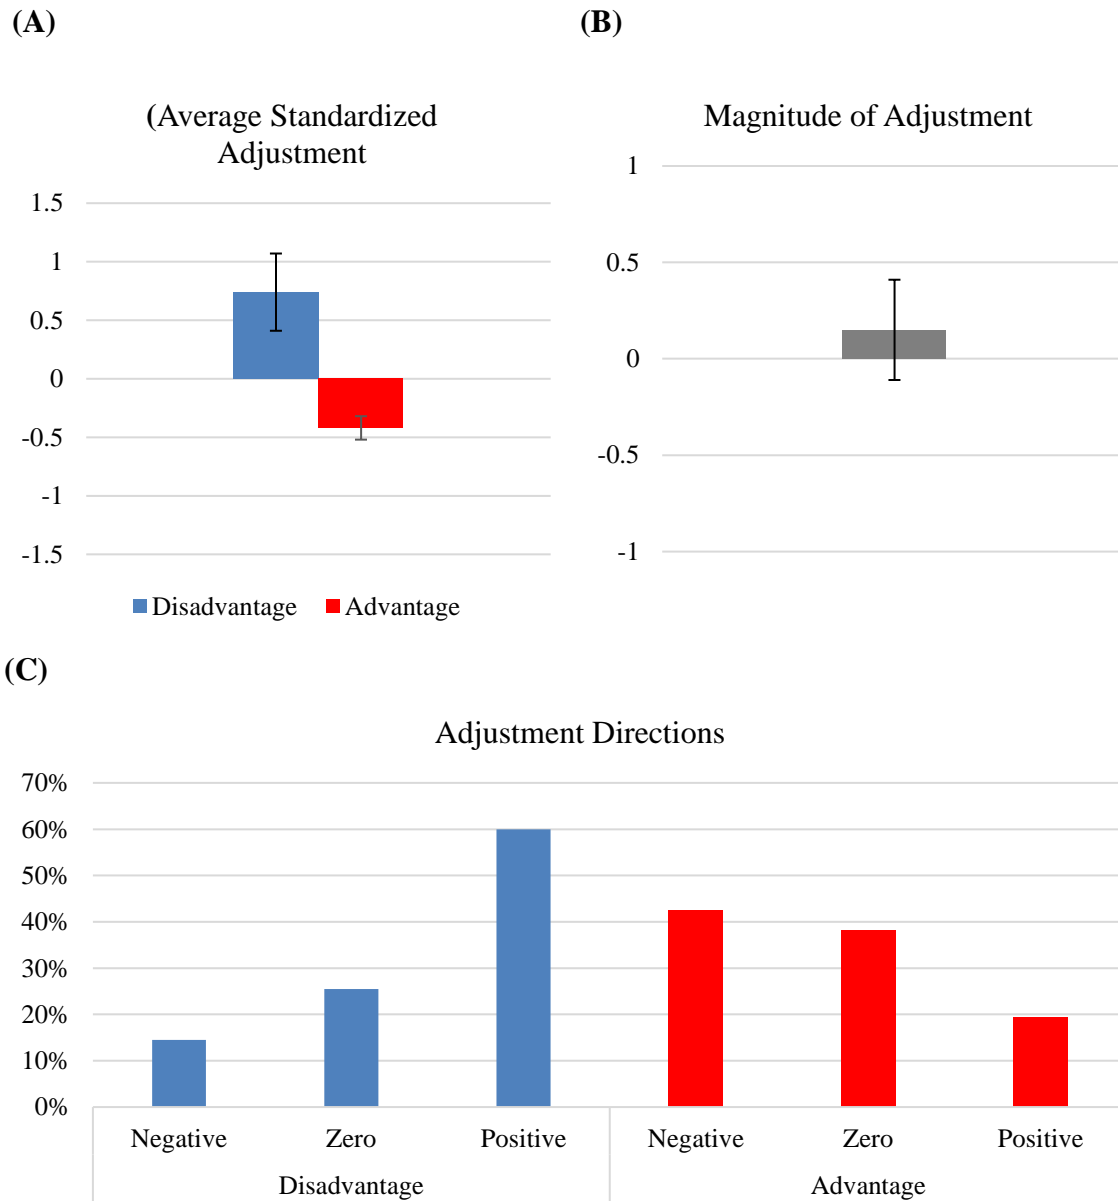
| | | |
|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------|-----------------------|
| 14/100 Overall Disadvantage Level <small>(based on the information below)</small> | | High Advantage |
| High School Information | | |
| School Type Public | Per pupil funding \$8,700 (83 rd percentile; high) | |
| Free/Reduced Lunch Eligibility 40% (low) | College Enrollment (3-yr avg.) 88% (high) | |
| Neighborhood Information | | |
| Median Household Income \$11,500 (75 th percentile nationwide) | | |
| Educational Attainment Index (based on high school and college completion rates) 5/5 (Very High) | | |
| Neighborhood Crime Index (based on frequency of 7 crime types) 1/5 (Very Low Crime) | | |

(B) Disadvantage

| | | |
|----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|--------------------------|
| 86/100 Overall Disadvantage Level <small>(based on the information below)</small> | | High Disadvantage |
| High School Information | | |
| School Type Public | Per pupil funding \$2,500 (17 th percentile; low) | |
| Free/Reduced Lunch Eligibility 96% (high) | College Enrollment 64% (low) | |
| Neighborhood Information | | |
| Median Household Income \$31,200 (25 th percentile nationwide) | | |
| Educational Attainment Index (based on high school and college completion rates) 1/5 (Very Low) | | |
| Neighborhood Crime Index (based on frequency of 7 crime types) 5/5 (Very High Crime) | | |

Note. Applicant’s socioeconomic **(A)** advantage and **(B)** disadvantage. The stimuli are based on the type of high school and neighborhood information that college admissions officers are able to access when evaluating applicants.

Figure 2.2 Results of Meta-Analysis, Chapter 2



Note. The meta-analysis examined adjustment for advantage versus disadvantage. **(A)** shows the average adjustment effect size and direction, **(B)** shows the effect size for the difference between advantage and disadvantage in absolute magnitude of adjustment, conditional on adjusting, and **(C)** shows the proportion of people who adjusted positively, negatively, or did not adjust.

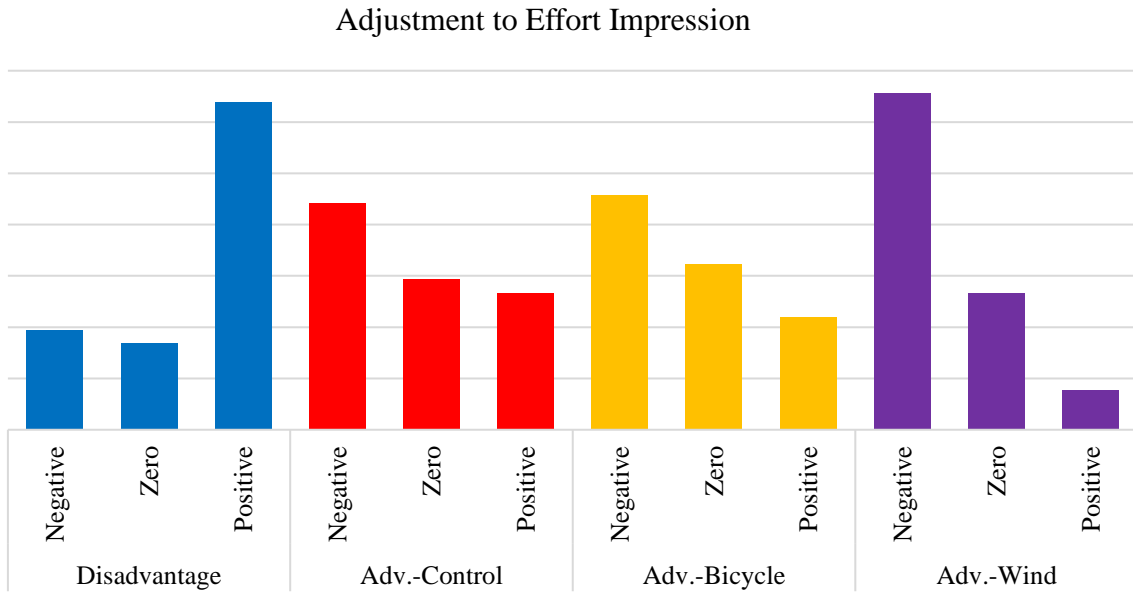
Figure 2.3 Effort Adjustment by Political Attitudes, Chapter 2

| | | | Political Attitudes | | | | | | |
|-------------------|--------------|----------|---------------------|---------|-----------|----------|------------|--------------|------------|
| | | | Ext. Lib. | Liberal | Sli. Lib. | Moderate | Sli. Cons. | Conservative | Ext. Cons. |
| Effort Adjustment | Advantage | Positive | 5.04% | 7.64% | 9.29% | 21.88% | 12.87% | 14.55% | 15.69% |
| | | Zero | 30.25% | 34.03% | 43.57% | 51.56% | 49.50% | 60.00% | 45.10% |
| | | Negative | 64.71% | 58.33% | 47.14% | 26.56% | 37.62% | 25.45% | 39.22% |
| | Disadvantage | Positive | 76.53% | 72.48% | 69.53% | 60.26% | 57.41% | 57.89% | 63.89% |
| | | Zero | 21.43% | 17.45% | 17.97% | 29.49% | 33.33% | 27.63% | 25.00% |
| | | Negative | 2.04% | 10.07% | 12.50% | 10.26% | 9.26% | 14.47% | 11.11% |

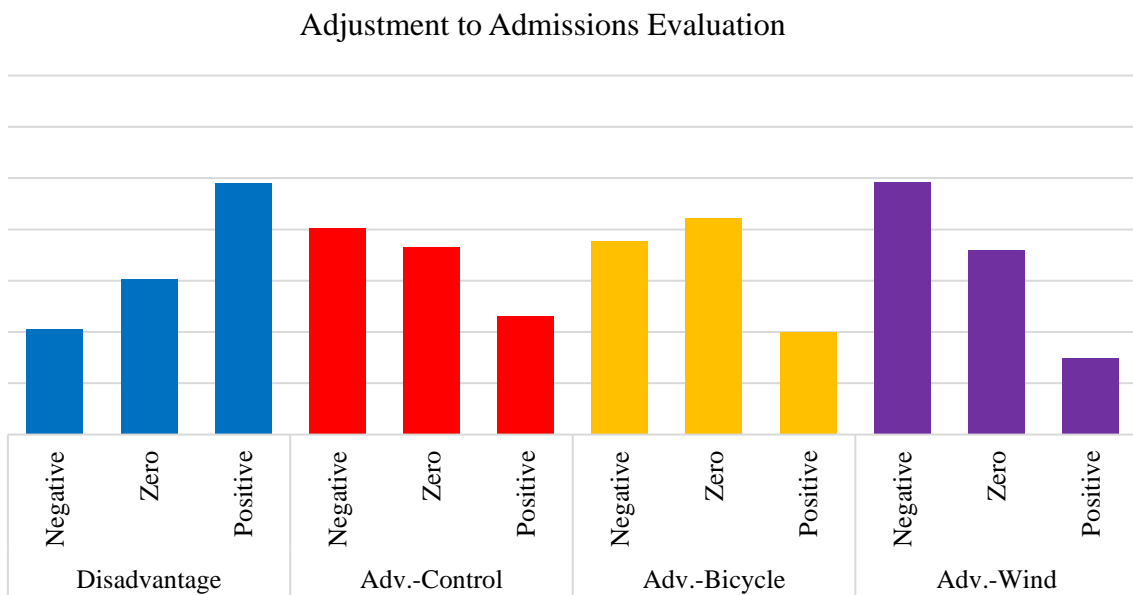
Note. Heat map of the most common effort adjustment (positive, zero, negative) when evaluating either an advantaged or disadvantaged applicant, split by self-reported political attitude.

Figure 2.4 Adjustment Directions, Chapter 2

(A)



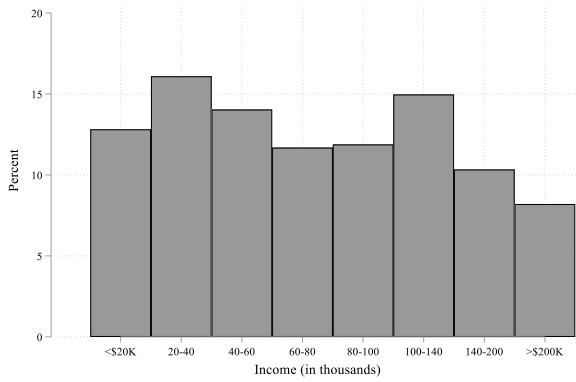
(B)




Note. Shows the proportion of people in each experimental condition who adjusted positively, negatively, or did not adjust their (A) effort impressions and (B) admissions evaluations.

Figure 2.5 Income Distributions, Chapter 2


(A) U.S. Income Distribution




(B) No Adjustment



(C) Symmetric Adjustment

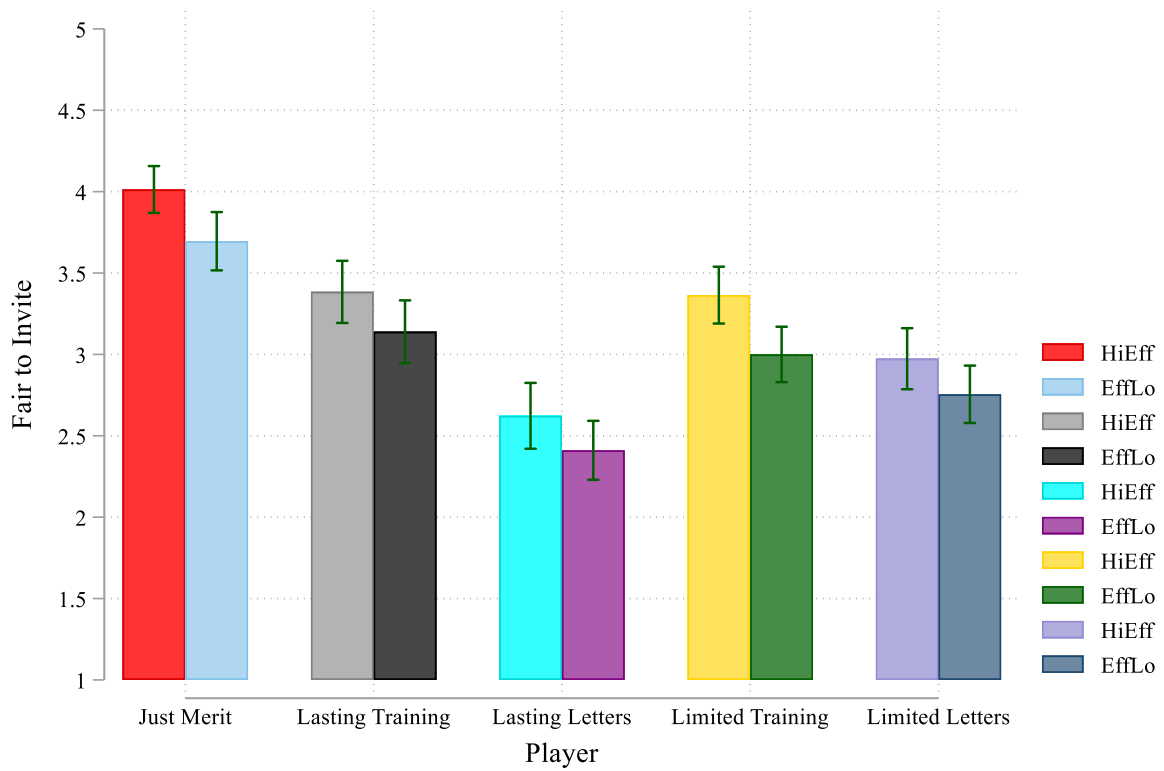


(D) Asymmetric Adjustment




Note. Income distributions of **(A)** the U.S. and **(B-D)** college classes admitted through different simulated evaluation processes: **(B)** no adjustment to SAT score percentile based on income, **(C)** symmetric adjustment: 2- to 8-percentile point adjustments based on distance from \$80,000 income, and **(D)** asymmetric adjustment: 2- to 8-percentile point adjustments for all incomes below \$80,000 and -2- to 8-percentile point adjustment for 54% of incomes above \$80,000.

Figure 3.1 Fairness Judgment, Chapter 3



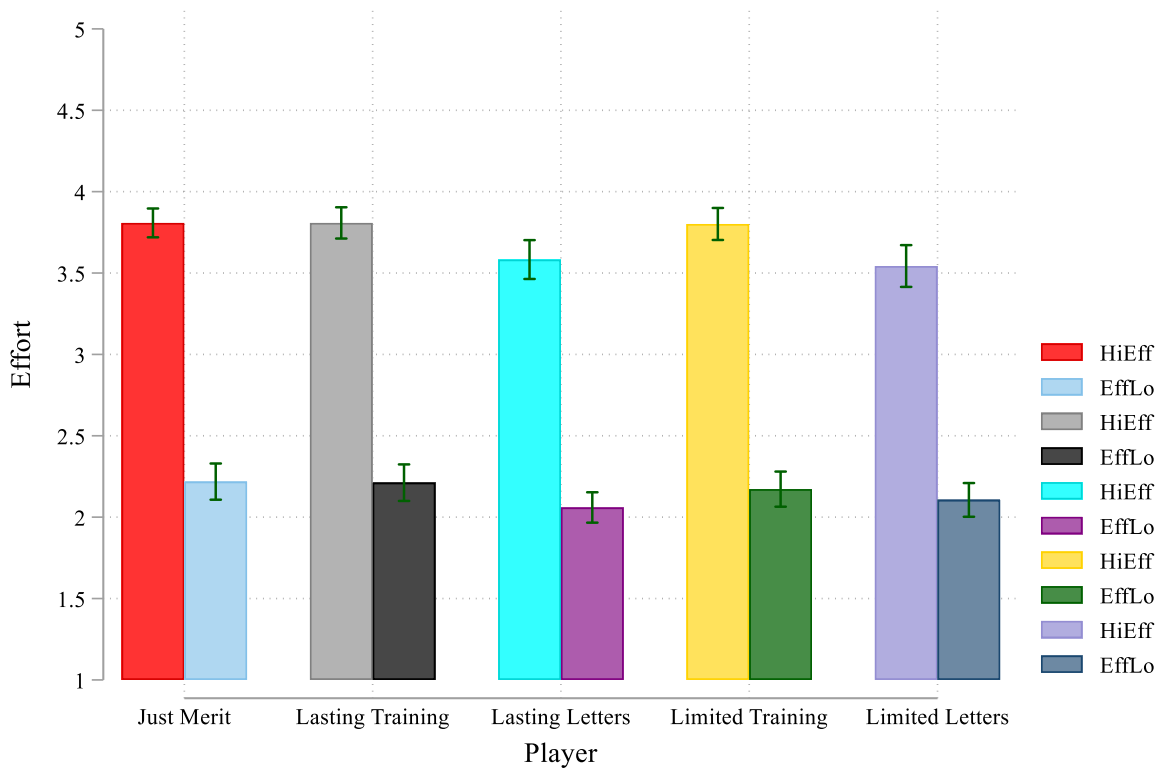
Note. Participants indicated whether it would unfair or fair to invite each Player to the tournament (1 = *Very Unfair*, 3 = *Neither*, 5 = *Very Fair*). Error bars show 95% confidence intervals.

Figure 3.2 Choice to Invite, Chapter 3




Note. Participants indicated whether they would invite each Player to the tournament if they were paid a bonus based on how well the Player performed in the tournament (1 = *Definitely Not*, 3 = *Maybe*, 5 = *Definitely Yes*). Error bars show 95% confidence intervals.

Figure 3.3 Effort, Chapter 3



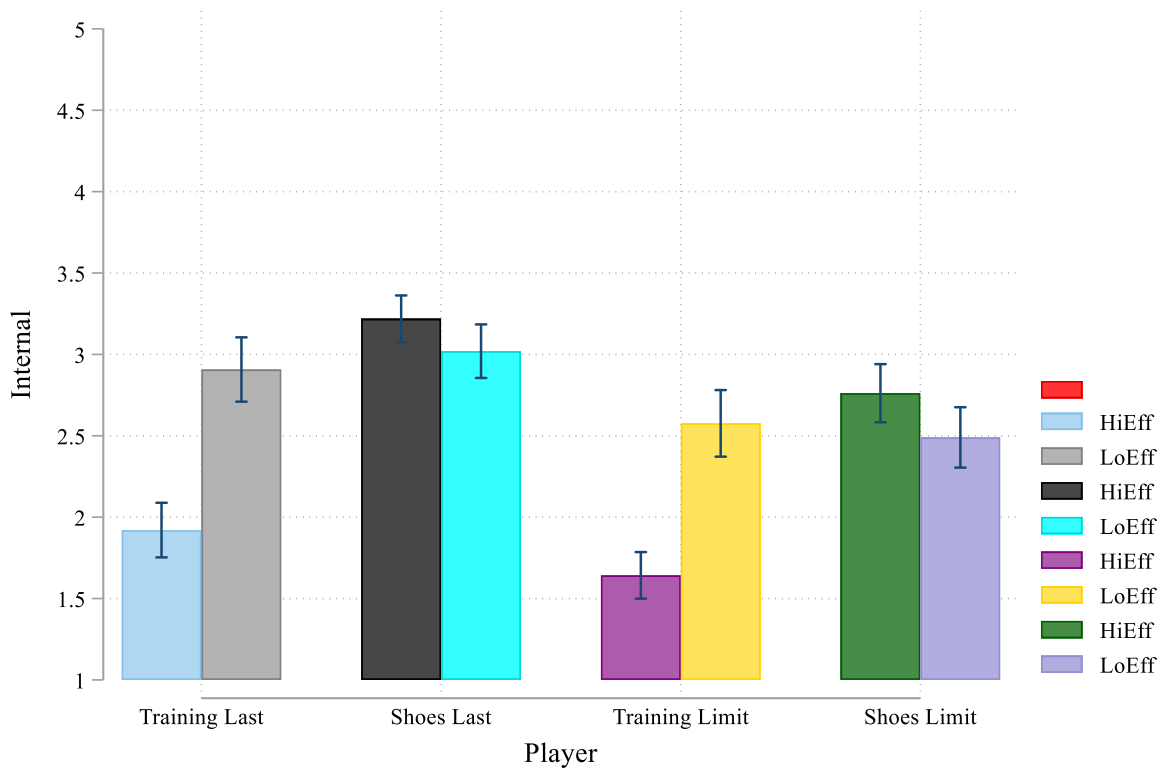
Note. Participants indicated how much effort the Player put into solving the Wordle puzzles (1=No effort at all to 5=An extreme amount of effort). Error bars show 95% confidence intervals.

Figure 3.4 Skill, Chapter 3



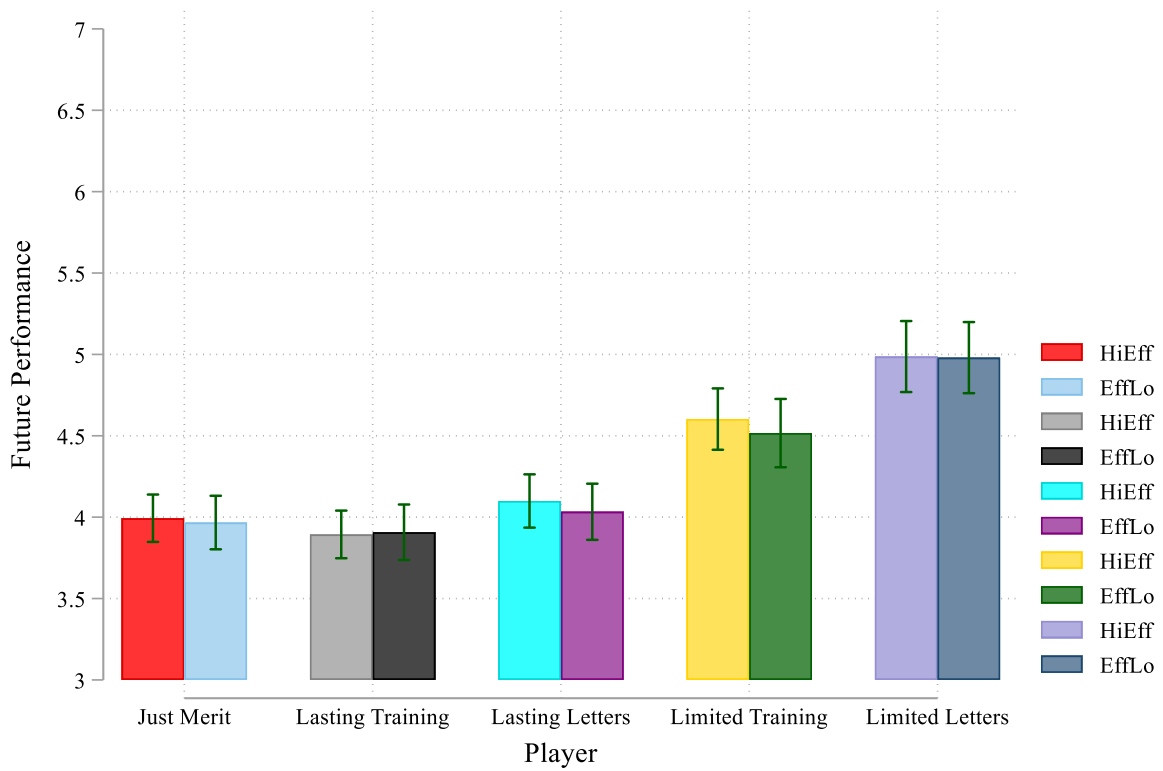
Note. Participants indicated how skilled they thought the Player was at Wordle puzzles (1=Not skilled at all to 5=Extremely skilled). Error bars show 95% confidence intervals.

Figure 3.5 Internal, Chapter 3



Note. Participants indicated whether the Player's situation actually improved their skill at Wordle (1=Did not improve at all to 5=Extremely improved). Error bars show 95% confidence intervals.

Figure 3.6 Tournament Prediction, Chapter 3



Note. Participants indicated how many guesses it would take the Player to solve the tournament puzzles, if they Player were to play in the tournament (1 guess, ..., 6 guesses, 7 = Will not be able to solve the puzzles). Error bars show 95% confidence intervals.

References

- Addona, V., & Yates, P. A. (2010). A closer look at the relative age effect in the National Hockey League. *Journal of Quantitative Analysis in Sports*, 6(4).
- Adler, P., & Jermier, J. (2005). Developing a field with more soul: Standpoint theory and public policy research for management scholars. *Academy of Management Journal*, 48(6), 941-944.
- Alesina, A., & Angeletos, G. M. (2005). Fairness and redistribution. *American Economic Review*, 95(4), 960-980.
- Alesina, A. F., Glaeser, E. L., & Sacerdote, B. (2001). *Why doesn't the US have a European-style welfare system?*. (No. 8524). National Bureau of Economic Research.
- Almås, I., Cappelen, A. W., Sørensen, E. Ø., & Tungodden, B. (2010). Fairness and the development of inequality acceptance. *Science*, 328(5982), 1176-1178.
- Alper, S. (2020). Explaining the complex effect of construal level on moral and political attitudes. *Current Directions in Psychological Science*, 29(2), 115-120.
- Ambrose, M. L., & Schminke, M. (2003). Organization structure as a moderator of the relationship between procedural justice, interactional justice, perceived organizational support, and supervisory trust. *Journal of Applied Psychology*, 88(2), 295 - 305.
- Arce, E., & Segura, D. A. (2015). Stratification in the labor market. *The Wiley Blackwell encyclopedia of race, ethnicity, and nationalism*, 1-3.
- Baker, D. J., & Bastedo, M. N. (2022). What If We Leave It Up to Chance? Admissions Lotteries and Equitable Access at Selective Colleges. *Educational Researcher*, 51(2), 134-145.
- Ballew, C. C., & Todorov, A. (2007). Predicting political elections from rapid and unreflective face judgments. *Proceedings of the National Academy of Sciences*, 104(46), 17948-17953.
- Ballinger, T., & Crocker, J. (2021). Understanding Whites' perceptions of multicultural policies: A (non) zero-sum framework?. *Journal of Personality and Social Psychology*, 120(5), 1231 - 1260.
- Bastedo, M. N., Bell, D. W., Howell, J. S., Hsu, J., Hurwitz, M., Perfetto, G., & Welch, M. (2022). Admitting students in context: Field experiments on information dashboards in college admissions. *The Journal of Higher Education*, 93(3), 327-374.
- Bastedo, M. N., Bowman, N. A., Glasener, K. M., & Kelly, J. L. (2018). What are we talking about when we talk about holistic review? Selective college admissions and its effects on low-SES students. *The Journal of Higher Education*, 89(5), 782-805.

- Bastedo, M. N., & Jaquette, O. (2011). Running in place: Low-income students and the dynamics of higher education stratification. *Educational Evaluation and Policy Analysis*, 33(3), 318-339.
- Baron, J. N., & Cook, K. S. (1992). Process and outcome: Perspectives on the distribution of rewards in organizations. *Administrative Science Quarterly*, 37(2), 191-197.
- Bazerman, M. H., & Tenbrunsel, A. E. (2011). Blind spots. In *Blind Spots*. Princeton University Press.
- Beamish, R. (1990). The Persistence of Inequality: An Analysis of Participation Patterns Among Canada's High Performance Athletes. *International Review for the Sociology of Sport*, 25(2), 143-155.
- Bell, J. M., & Hartmann, D. (2007). Diversity in everyday discourse: The cultural ambiguities and consequences of “happy talk”. *American Sociological Review*, 72(6), 895-914.
- Benabou, R., & Tirole, J. (2006). Belief in a just world and redistributive politics. *The Quarterly Journal of Economics*, 121(2), 699-746.
- Benartzi, S., Beshears, J., Milkman, K. L., Sunstein, C. R., Thaler, R. H., Shankar, M., ... & Galing, S. (2017). Should governments invest more in nudging?. *Psychological Science*, 28(8), 1041-1055.
- 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.
- Blanchard, T. C. J. “Resolution of compliance review of Princeton University’s undergraduate admissions” (Case No. 02-08-6002, U.S. Department of Education Office for Civil Rights, 2015); <https://www2.ed.gov/about/offices/list/ocr/docs/investigations/02086002-a.pdf>
- Bobocel, D. R., Son Hing, L. S., Holmvall, C. M., & Zanna, M. P. (2002). Policies to redress social injustice: Is the concern for justice a cause both of support and of opposition? In M. Ross & D. Miller (Eds.), *The Justice Motive in Everyday Life*: 204–225. New York, NY: Cambridge University Press.
- Bohnet, I., Van Geen, A., & Bazerman, M. (2016). When performance trumps gender bias: Joint vs. separate evaluation. *Management Science*, 62(5), 1225-1234.
- Borowiecki, K. (2019). The origins of creativity: The case of the arts in the United States since 1850. *Discussion Papers on Business and Economics*, University of Southern Denmark, 3.

- Bowles, S., Gintis, H., & Osborne Groves, M. (Eds.). (2009). *Unequal chances: Family background and economic success*. Princeton University Press.
- Buhrmester, M. D., Talaifar, S., & Gosling, S. D. (2018). An evaluation of Amazon's Mechanical Turk, its rapid rise, and its effective use. *Perspectives on Psychological Science*, *13*(2), 149-154.
- Bullock, H. E., Williams, W. R., & Limbert, W. M. (2003). Predicting support for welfare policies: The impact of attributions and beliefs about inequality. *Journal of Poverty*, *7*(3), 35-56.
- Brickman, P., Folger, R., Goode, E., & Schul, Y. (1981). Microjustice and macrojustice. In M. J. Lerner & S. C. Lerner (Eds.), *The justice motive in social behavior* (pp. 173–204). New York, NY: Plenum.
- Cappelen, A. W., Eichele, T., Hugdahl, K., Specht, K., Sørensen, E. Ø., & Tungodden, B. (2014). Equity theory and fair inequality: A neuroeconomic study. *Proceedings of the National Academy of Sciences*, *111*(43), 15368-15372.
- Cappelen, A. W., Konow, J., Sørensen, E. Ø., & Tungodden, B. (2013). Just luck: An experimental study of risk-taking and fairness. *American Economic Review*, *103*(4), 1398-1413.
- Cappelen, A. W., Mollerstrom, J., Reme, B. A., & Tungodden, B. (2022). A meritocratic origin of egalitarian behaviour. *The Economic Journal*, *132*(646), 2101-2117.
- Cappelen, A. W., Moene, K. O., Skjelbred, S. E., & Tungodden, B. (2023). The merit primacy effect. *The Economic Journal*, *133*(651), 951-970.
- Carnevale, A. P., & Strohl, J. (2010). "How increasing college access is increasing inequality, and what to do about it." In R. D. Kahlenberg (Ed.), *Rewarding strivers: Helping low-income students succeed in college* (pp. 71-190). Century Foundation Press.
- CBS News & New York Times. (1997, December). CBS News/New York Times poll # 1997-97012A: Campaign fundraising/parents/affirmative action.
- Chang, E. H., Kirgios, E. L., Rai, A., & Milkman, K. L. (2020). The isolated choice effect and its implications for gender diversity in organizations. *Management Science*, *66*(6), 2752-2761.
- Cheryan, S., Ziegler, S. A., Montoya, A. K., & Jiang, L. (2017). Why are some STEM fields more gender balanced than others?. *Psychological bulletin*, *143*(1), 1 – 35.
- Chetty, R., Friedman, J. N., Saez, E., Turner, N., & Yagan, D. (2020). Income segregation and intergenerational mobility across colleges in the United States. *The Quarterly Journal of Economics*, *135*(3), 1567-1633.

- Chetty, R., Friedman, J. N., Saez, E., Turner, N., & Yagan, D. (2017). *Mobility report cards: The role of colleges in intergenerational mobility* (No. w23618). National Bureau of Economic Research.
- Chetty, R., Hendren, N., Kline, P., & Saez, E. (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *The Quarterly Journal of Economics*, 129(4), 1553-1623.
- Clinedinst, M. (2020). "State of College Admission 2019" *National Association for College Admission Counseling*.
- Colby, H., DeWitt, J., & Chapman, G. B. (2015). Grouping promotes equality: the effect of recipient grouping on allocation of limited medical resources. *Psychological Science*, 26(7), 1084-1089.
- College Board (2016). "2016 Total Group SAT Suite of Assessments Annual Report". <https://reports.collegeboard.org/media/pdf/2016-total-group-sat-suite-assessments-annual-report.pdf>
- College Board (2016). "Insight into nonacademic factors and practice" (Report on Future Admissions Tools and Models Initiative). <https://secure-media.collegeboard.org/digitalServices/pdf/professionals/insight-into-nonacademic-factors-and-practice.pdf>
- Colquitt, J. A. (2001). On the dimensionality of organizational justice: a construct validation of a measure. *Journal of Applied Psychology*, 86(3), 386 - 400.
- Colquitt, J. A., Conlon, D. E., Wesson, M. J., Porter, C. O., & Ng, K. Y. (2001). Justice at the millennium: a meta-analytic review of 25 years of organizational justice research. *Journal of Applied Psychology*, 86(3), 425 - 445.
- Colquitt, J. A., LePine, J. A., Piccolo, R. F., Zapata, C. P., & Rich, B. L. (2012). Explaining the justice–performance relationship: Trust as exchange deepener or trust as uncertainty reducer?. *Journal of Applied Psychology*, 97(1), 1-15.
- Colquitt, J. A., Sabey, T. B., Rodell, J. B., & Hill, E. T. (2019). Content validation guidelines: Evaluation criteria for definitional correspondence and definitional distinctiveness. *Journal of Applied Psychology*, 104(10), 1243-1265.
- Colquitt, J. A., Scott, B. A., Rodell, J. B., Long, D. M., Zapata, C. P., Conlon, D. E., & Wesson, M. J. (2013). Justice at the millennium, a decade later: a meta-analytic test of social exchange and affect-based perspectives. *Journal of Applied Psychology*, 98(2), 199-236.
- Critcher, C. R., & Dunning, D. (2013). Predicting persons' versus a person's goodness: Behavioral forecasts diverge for individuals versus populations. *Journal of Personality and Social Psychology*, 104(1), 28-44.

- Crosby, F. J., Iyer, A., Clayton, S., & Downing, R. A. (2003). Affirmative action: Psychological data and the policy debates. *American Psychologist*, 58(2), 93.
- Crosby, F. J., Iyer, A., & Sincharoen, S. (2006). Understanding affirmative action. *Annual Review of Psychology*, 57, 585-611.
- Crowne, D. P., & Marlowe, D. (1960). A new scale of social desirability independent of psychopathology. *Journal of Consulting Psychology*, 24(4), 349-354.
- Davidai, S. (2022). How do people make sense of wealth and poverty?. *Current Opinion in Psychology*, 43, 42-47.
- Davidai, S., & Gilovich, T. (2016). The headwinds/tailwinds asymmetry: An availability bias in assessments of barriers and blessings. *Journal of Personality and Social Psychology*, 111(6), 835.
- Davidai, S., & Gilovich, T. (2018). How should we think about Americans' beliefs about economic mobility?. *Judgment and Decision Making*, 13(3), 297-304.
- Davidai, S., & Ongis, M. (2019). The politics of zero-sum thinking: The relationship between political ideology and the belief that life is a zero-sum game. *Science Advances*, 5(12), eaay3761.
- 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.
- Dixon, J., Durrheim, K., & Tredoux, C. (2007). Intergroup contact and attitudes toward the principle and practice of racial equality. *Psychological Science*, 18(10), 867-872.
- Donovan, S.A., Dalaker, J., Labonte, M., Romero, P. D. (2021). The U.S. Income Distribution: Trends and Issues (Rep. R44705, Congressional Research Service). <https://fas.org/sgp/crs/misc/R44705.pdf>
- Durante, R., Putterman, L., & Van der Weele, J. (2014). Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4), 1059-1086.
- Einav, L., & Yariv, L. (2006). What's in a surname? The effects of surname initials on academic success. *Journal of Economic Perspectives*, 20(1), 175-188.
- Evers, E. R., Inbar, Y., Blanken, I., & Oosterwijk, L. D. (2017). When do people prefer carrots to sticks? A robust "matching effect" in policy evaluation. *Management Science*, 63(12), 4261-4276.
- Fritz, M. S., & MacKinnon, D. P. (2007). Required sample size to detect the mediated effect. *Psychological Science*, 18(3), 233-239.

- Gasper, K., & Clore, G. L. (2002). Attending to the big picture: Mood and global versus local processing of visual information. *Psychological science, 13*(1), 34-40.
- Goldin, C., & Rouse, C. (2000). Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American Economic Review, 90*(4), 715-741.
- Gomila, R. (2021). Logistic or linear? Estimating causal effects of experimental treatments on binary outcomes using regression analysis. *Journal of Experimental Psychology: General, 150*(4), 700-709.
- Harrison, D. A., Kravitz, D. A., Mayer, D. M., Leslie, L. M., & Lev-Arey, D. (2006). Understanding attitudes toward affirmative action programs in employment: Summary and meta-analysis of 35 years of research. *Journal of Applied Psychology, 91*(5), 1013-1036.
- Hayes, A. F. (2017). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. Guilford publications.
- Hayes, A. F. 2018. *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (2nd ed.). New York, NY: Guilford Publications.
- Hideg, I., & Ferris, D. L. (2017). Dialectical thinking and fairness-based perspectives of affirmative action. *Journal of Applied Psychology, 102*(5), 782-801.
- Hideg, I., Michela, J. L., & Ferris, D. L. (2011). Overcoming negative reactions of nonbeneficiaries to employment equity: The effect of participation in policy formulation. *Journal of Applied Psychology, 96*(2), 363-376.
- Hinkin, T. R., & Tracey, J. B. (1999). An analysis of variance approach to content validation. *Organizational Research Methods, 2*(2), 175-186.
- Ho, G. C., & Unzueta, M. M. (2015). Antiegalitarians for affirmative action? When social dominance orientation is positively related to support for egalitarian social policies. *Journal of Applied Social Psychology, 45*(8), 451-460.
- Hoxby, C. M., & Turner, S. (2015). What high-achieving low-income students know about college. *American Economic Review, 105*(5), 514-517.
- Humphreys, A., Isaac, M. S., & Wang, R. J. H. (2021). Construal matching in online search: applying text analysis to illuminate the consumer decision journey. *Journal of Marketing Research, 58*(6), 1101-1119.
- Huppert, E., Shaw, A., & Decety, J. (2020). The effect of hunger on children’s sharing behavior and fairness preferences. *Journal of Experimental Child Psychology, 192*, 104786.
- Jones, E. E. (1990). *Interpersonal perception*. WH Freeman/Times Books/Henry Holt & Co.

- Jost, J. T., Federico, C. M., & Napier, J. L. (2009). Political ideology: Its structure, functions, and elective affinities. *Annual Review of Psychology, 60*, 307-337.
- Kteily, N. S., Sheehy-Skeffington, J., & Ho, A. K. (2017). Hierarchy in the eye of the beholder:(Anti-) egalitarianism shapes perceived levels of social inequality. *Journal of Personality and Social Psychology, 112*(1), 136-159.
- Kraus, M. W., Torrez, B., Park, J. W., & Ghayebi, F. (2019). Evidence for the reproduction of social class in brief speech. *Proceedings of the National Academy of Sciences, 116*(46), 22998-23003.
- Kraus, M. W., & Tan, J. J. (2015). Americans overestimate social class mobility. *Journal of Experimental Social Psychology, 58*, 101-111.
- Krawczyk, M. (2010). A glimpse through the veil of ignorance: Equality of opportunity and support for redistribution. *Journal of Public Economics, 94*(1-2), 131-141.
- Laham, S. M., Koval, P., & Alter, A. L. (2012). The name-pronunciation effect: Why people like Mr. Smith more than Mr. Colquhoun. *Journal of Experimental Social Psychology, 48*(3), 752-756.
- Liberman, N., & Trope, Y. (1998). The role of feasibility and desirability considerations in near and distant future decisions: A test of temporal construal theory. *Journal of personality and social psychology, 75*(1), 5-18.
- Lowery, B. S., Unzueta, M. M., Knowles, E. D., & Goff, P. A. (2006). Concern for the in-group and opposition to affirmative action. *Journal of Personality and Social Psychology, 90*(6), 961-974.
- Lucas, B. J., Berry, Z., Giurge, L. M., & Chugh, D. (2021). A longer shortlist increases the consideration of female candidates in male-dominant domains. *Nature Human Behaviour, 5*(6), 736-742.
- Luguri, J. B., Napier, J. L., & Dovidio, J. F. (2012). Reconstructing intolerance: Abstract thinking reduces conservatives' prejudice against nonnormative groups. *Psychological Science, 23*(7), 756-763.
- Mahfud, Y., Badea, C., Verkuyten, M., & Reynolds, K. (2018). Multiculturalism and attitudes toward immigrants: The impact of perceived cultural distance. *Journal of Cross-Cultural Psychology, 49*(6), 945-958.
- McFarlin, D. B., & Sweeney, P. D. (1992). Distributive and procedural justice as predictors of satisfaction with personal and organizational outcomes. *Academy of Management Journal, 35*(3), 626-637.

- Mentovich, A., Yudkin, D., Tyler, T., & Trope, Y. (2016). Justice without borders: The influence of psychological distance and construal level on moral exclusion. *Personality and Social Psychology Bulletin*, *42*(10), 1349-1363.
- Milanovic, B. (2015). Global inequality of opportunity: How much of our income is determined by where we live?. *Review of Economics and Statistics*, *97*(2), 452-460.
- Milkman, K. L., Akinola, M., & Chugh, D. (2015). What happens before? A field experiment exploring how pay and representation differentially shape bias on the pathway into organizations. *Journal of Applied Psychology*, *100*(6), 1678-1712.
- Milkman, K. L., Mazza, M. C., Shu, L. L., Tsay, C. J., & Bazerman, M. H. (2012). Policy bundling to overcome loss aversion: A method for improving legislative outcomes. *Organizational Behavior and Human Decision Processes*, *117*(1), 158-167.
- Milkman, K. L., Rogers, T., & Bazerman, M. H. (2008). Harnessing our inner angels and demons: What we have learned about want/should conflicts and how that knowledge can help us reduce short-sighted decision making. *Perspectives on Psychological Science*, *3*(4), 324-338.
- Nakashima, N. A., Daniels, D. P., & Laurin, K. (2017). It's about time: Divergent evaluations of restrictive policies in the near and distant future. *Organizational Behavior and Human Decision Processes*, *142*, 12-27.
- Napier, J. L., & Luguri, J. B. (2013). Moral mind-sets: Abstract thinking increases a preference for "individualizing" over "binding" moral foundations. *Social Psychological and Personality Science*, *4*(6), 754-759.
- NBC News & The Wall Street Journal. (1995, July). NBC News/Wall Street Journal poll: Clinton and federal budget.
- NBC News & The Wall Street Journal. (1997, December). NBC News/Wall Street Journal poll: Economy/American society.
- O'Laughlin, M. J., & Malle, B. F. (2002). How people explain actions performed by groups and individuals. *Journal of Personality and Social Psychology*, *82*(1), 33-48.
- Paulhus, D. L. (1991). Measurement and control of response bias. In J. P. Robinson, P. R. Shaver, & L. S. Wrightsman (Eds.), *Measures of personality and social psychological attitudes* (17-59). Academic Press.
- Pearce, J. L., & Huang, L. (2012). The decreasing value of our research to management education. *Academy of Management Learning & Education*, *11*(2), 247-262.

- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology, 70*, 153-163.
- 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.
- Phillips, L. T., Stephens, N. M., Townsend, S. S., & 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.
- Quadlin, N. (2018). The mark of a woman's record: Gender and academic performance in hiring. *American Sociological Review, 83*(2), 331-360.
- Rai, T. S., & Fiske, A. P. (2011). Moral psychology is relationship regulation: Moral motives for unity, hierarchy, equality, and proportionality. *Psychological Review, 118*(1), 57-75.
- Rattan, A., Savani, K., Chugh, D., & Dweck, C. S. (2015). Leveraging mindsets to promote academic achievement: Policy recommendations. *Perspectives on Psychological Science, 10*(6), 721-726.
- Read, D., & Loewenstein, G. (1995). Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. *Journal of Experimental Psychology: Applied, 1*(1), 34-49.
- Reardon, S. F. (2016). School district socioeconomic status, race, and academic achievement. *Stanford Center for Educational Policy Analysis*.
- Reeder, G. D., Hesson-McInnis, M., Krohse, J. O., & Scialabba, E. A. (2001). Inferences about effort and ability. *Personality and Social Psychology Bulletin, 27*(9), 1225-1235.
- Reyt, J. N., & Wiesenfeld, B. M. (2015). Seeing the forest for the trees: Exploratory learning, mobile technology, and knowledge workers' role integration behaviors. *Academy of Management Journal, 58*(3), 739-762.
- Rigol, G. W. & College Board (2003). *Admissions Decision-Making Models: How US Institutions of Higher Education Select Undergraduate Students*. College Board. <https://files.eric.ed.gov/fulltext/ED562589.pdf>
- Ritov, I., & Zamir, E. (2014). Affirmative action and other group tradeoff policies: Identifiability of those adversely affected. *Organizational Behavior and Human Decision Processes, 125*(1), 50-60.
- Rivera, L. A. (2011). Ivies, extracurriculars, and exclusion: Elite employers' use of educational credentials. *Research in Social Stratification and Mobility, 29*(1), 71-90.

- Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review*, *77*(6), 999-1022.
- Rivera, L. A., & Stevens, M. L. (2013). Why economic sociologists should care about education. *Accounts: ASA Economic Sociology Newsletter*, *12*(3), 5-8.
- Rivera, L. A., & Tilcsik, A. (2016). Class advantage, commitment penalty: The gendered effect of social class signals in an elite labor market. *American Sociological Review*, *81*(6), 1097-1131.
- Rogers, T., & Bazerman, M. H. (2008). Future lock-in: Future implementation increases selection of 'should' choices. *Organizational Behavior and Human Decision Processes*, *106*(1), 1-20.
- Sackett, P. R., Kuncel, N. R., Arneson, J. J., Cooper, S. R., & Waters, S. D. (2009). Does socioeconomic status explain the relationship between admissions tests and post-secondary academic performance?. *Psychological Bulletin*, *135*(1), 1-22.
- Smith, C. W., & Mayorga-Gallo, S. (2017). The new principle-policy gap: How diversity ideology subverts diversity initiatives. *Sociological Perspectives*, *60*(5), 889-911.
- Smith, P. K., & Trope, Y. (2006). You focus on the forest when you're in charge of the trees: power priming and abstract information processing. *Journal of Personality and Social Psychology*, *90*(4), 578-596.
- Soderberg, C. K., Callahan, S. P., Kochersberger, A. O., Amit, E., & Ledgerwood, A. (2015). The effects of psychological distance on abstraction: Two meta-analyses. *Psychological Bulletin*, *141*(3), 525-548.
- Soll, J. B., Keeney, R. L., & Larrick, R. P. (2013). Consumer misunderstanding of credit card use, payments, and debt: Causes and solutions. *Journal of Public Policy and Marketing*, *32*(1), 66-81.
- Starmans, C., Sheskin, M., & Bloom, P. (2017). Why people prefer unequal societies. *Nature Human Behaviour*, *1*(4), 1-7.
- Starck, J. G., Sinclair, S., & Shelton, J. N. (2021). How university diversity rationales inform student preferences and outcomes. *Proceedings of the National Academy of Sciences*, *118*(16), e2013833118.
- Stephens, N. M., Markus, H. R., & Phillips, L. T. (2014). Social class culture cycles: How three gateway contexts shape selves and fuel inequality. *Annual Review of Psychology*, *65*, 611-634.
- Sunshine, J., & Tyler, T. R. (2003). The role of procedural justice and legitimacy in shaping public support for policing. *Law and Society Review*, *37*(3), 513-548.

- Sukumar, P. T., & Metoyer, R. (2018). A visualization approach to addressing reviewer bias in holistic college admissions. *Cognitive Biases in Visualizations*, 161-175.
- Sukumar, P. T., Metoyer, R., & He, S. (2018). Making a pecan pie: Understanding and supporting the holistic review process in admissions. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), 1-22.
- Tannenbaum, D., Fox, C. R., & Rogers, T. (2017). On the misplaced politics of behavioural policy interventions. *Nature Human Behaviour*, 1(7), 0130.
- Tinghög, G., Andersson, D., & Västfjäll, D. (2017). Are individuals luck egalitarians?—An experiment on the influence of brute and option luck on social preferences. *Frontiers in Psychology*, 8, 460.
- Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728), 1623-1626.
- Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review*, 110(3), 403-421.
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review*, 117(2), 440-463.
- Tuch, S. A., & Hughes, M. (1996). Whites' racial policy attitudes. *Social Science Quarterly*, 723-745.
- Ubel, P. A., & Loewenstein, G. (1996). Distributing scarce livers: the moral reasoning of the general public. *Social Science and Medicine*, 42(7), 1049-1055.
- Van Meter, D. S., & Van Horn, C. E. (1975). The policy implementation process: A conceptual framework. *Administration and Society*, 6(4), 445-488.
- Wagstaff, G. F., Huggins, J. P., & Perfect, T. J. (1993). Equity, equality, and need in the adult family. *The Journal of Social Psychology*, 133(4), 439-443.
- Wiesenfeld, B. M., Reyt, J. N., Brockner, J., & Trope, Y. (2017). Construal level theory in organizational research. *Annual Review of Organizational Psychology and Organizational Behavior*, 4, 367-400.
- Yogeeswaran, K., & Dasgupta, N. (2014). The devil is in the details: Abstract versus concrete construals of multiculturalism differentially impact intergroup relations. *Journal of Personality and Social Psychology*, 106(5), 772-789.

Zdaniuk, A., & Bobocel, D. R. (2011). Independent self-construal and opposition to affirmative action: The role of microjustice and macrojustice preferences. *Social Justice Research, 24*, 341-364.

Zhang, F. "SAT Score Percentiles (High-Precision Version)" (PrepScholar, 2020); <https://blog.prepscholar.com/sat-percentiles-high-precision-2016>

Note on Supplementary Materials

The supplementary materials comprise:

- Additional analyses for the studies reported in Chapter 1
- Additional studies (summarized in Table 1.2) relevant to the findings in Chapter 1
- Details on the methods and results of the studies reported in Chapter 2

These materials are accessible in the following Open Science Framework (OSF) repository:

Munguia Gomez, D. M. (2023). Dissertation: Fairness and merit in the selection of advantaged and disadvantaged applicants. Retrieved from osf.io/aqs5z