SOCIUS
Original Article

Socius: Sociological Research for a Dynamic World
Volume 8: 1–12
© The Author(s) 2022
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/23780231221117888
srd.sagepub.com
SAGE

Insidious Nonetheless: How Small Effects and Hierarchical Norms Create and Maintain Gender Disparities in Organizations

Yuhao Du[1], Jessica Nordell[2], and Kenneth Joseph[1]

Abstract
The term *glass ceiling* is applied to the well-established phenomenon in which women and people of color are consistently blocked from reaching the uppermost levels of the corporate hierarchy. Focusing on gender, we present an agent-based model that explores how empirically established mechanisms of interpersonal discrimination coevolve with social norms at both the organizational (meso) and societal (macro) levels to produce this glass ceiling effect for women. Our model extends the understanding of how the glass ceiling arises and why it can be resistant to change. We do so by synthesizing existing psychological and structural theories of discrimination into a mathematical model that quantifies explicitly how complex organizational systems can produce and maintain inequality. We discuss implications of our findings for both intervention and future empirical analyses and provide open-source code for those wishing to adapt or extend our work.

Keywords
agent-based modeling, social norms, gender inequality

Men are overrepresented at higher levels of the corporate hierarchy. *The New York Times* reports, for instance, that in 2018, there were fewer female chief executives at *Fortune* 500 companies than male chief executives with the name James, despite the fact that only 3.3 percent of the U.S. population is named James while women make up 50.3 percent of the U.S. population (Miller, Quealy, and Sanger-Katz 2018).

Scholars have long studied potential reasons for this *glass ceiling effect*—the name given for the general phenomenon in which invisible barriers block women and people of color from reaching high levels of management (Bertrand 2018; Cotter et al. 2001). Here, we focus on literature specifically surrounding gender. Some popular explanations of the glass ceiling revolve around innate or learned differences between men and women, such as psychological differences in risk-taking or taste for competition/negotiation (Babcock and Laschever 2009; Reuben, Sapienza, and Zingales 2015; Schubert, Brown, and Brachinger 2000), or differences on personality traits (Collischon 2021; Filer 1983; Semykina and Linz 2007). Others have focused beyond the individual, to the places where gender norms—roughly, culturally prescribed guidelines for behavior based on one's own perceived gender and the perceived gender of those around us—and stereotypes—generalized and often unfounded assumptions about someone based on their (perceived) gender—are learned and enforced. To this end, scholars have found that policy, including family leave and flexible scheduling (Bear 2021; Goldin and Katz 2016; Pettit and Hook 2005; Williams and Segal 2003), and interpersonal factors such as harassment (Berdahl and Moore 2006; Stockdale and Bhattacharya 2009) and gender-biased evaluations (Heilman and Haynes 2005; Moss-Racusin et al. 2012), both play significant roles in creating or limiting the upward mobility of women in the workplace.

Due to the limits of what can be operationalized in a single study, efforts to empirically identify causes of the glass ceiling rarely consider more than a few competing ideas and often do so at a single (or a few) moments in time. This can be problematic because corporations are examples of

[1]University at Buffalo, Buffalo, NY, USA
[2]Unaffiliated, Minneapolis, MN, USA

Corresponding Author:
Kenneth Joseph, Computer Science and Engineering Department, University at Buffalo, 335 Davis Hall, Buffalo, NY, USA.
Email: kjoseph@buffalo.edu

complex social systems (Harrison et al. 2007; Martell, Emrich, and Robison-Cox 2012), where social norms and stereotypes diffuse over time through individuals and groups within the organization and back and forth between the organization and society. The interaction of these multiple and hierarchical social structures creates feedback processes and unexpected outcomes that can confound simple explanations of empirical findings. Empirical work therefore cannot always cleanly capture or quantify the ways in which women experience gender discrimination in many ways over long periods of time.

Acknowledging the limitations of empirical work to understand complex social systems, scholars have turned to *simulation*, and in particular, *agent-based modeling*, to study gender disparities in organizations (Bullinaria 2018; Martell, Lane and Emrich 1996; Momennejad, Sinclair, and Cikara 2019; Robison-Cox, Martell, and Emrich 2007). In an agent-based model, a computational, simplified representation of an individual (an "agent") interacts with other agents using a predefined set of rules. These rules shape macro-level statistics, which can then "feed back" to reshape the parameters of the established rules (Gilbert 2007). Agent-based models have long been used in the social sciences to study phenomena within complex systems because one can rapidly consider experiments that are too large for empirical study and can also easily examine counterfactual arguments within evolving systems (Carley 1991).

The present work proposes a new agent-based model of how the glass ceiling emerges within the complex social system of a hypothetical corporation. We outline how glass ceilings within organizations can emerge through a coupling of (1) stable, hierarchical gendered norms about whose contributions are valued and how and (2) small, discrete instances in which these norms are enacted at the interpersonal level.[1] We use this new model of the glass ceiling effect to study how, together, these impact the success or failure of a quota-based intervention.

Model Overview

Our work is based on the earlier simulation model of Martell et al. (1996), who show how gender disparities in the corporate hierarchy can arise solely from small gender biases during performance evaluation. As in their work, agents in our model represent employees of a hypothetical, eight-level corporate organization, with a prespecified number of agents at each level. The primary difference between our model and Martell et al. (1996) is that we link gender disparity not to a generalized notion of "bias" in performance evaluations, but to specific, empirically identified mechanisms through which this bias manifests. To do so, our model simulates two common process in organizations: Employees engage in projects, and employees are promoted through the ranks of the company. Projects may succeed or fail, and promotions are based on the agents' *perceived promotability*.

At the start of the simulation, agents are randomly initialized with a perceived binary gender (man or woman) and a perceived promotability. The simulation then iterates over a series of *turns*. On each turn, agents receive either an individual or group *project*. At fixed intervals, we also introduce stretch projects that provide outsized boosts in perceived promotability. Projects randomly succeed or fail with equal probability. When an agent's project succeeds, the agent receives some *credit* that increases their perceived promotability. When a project fails, the agent's perceived promotability drops via some amount of credit.

After some number of simulation turns, there is a *promotion cycle*. During a promotion cycle, the employees with the highest perceived promotability move up from their current level of the corporate hierarchy to the next. In the rare case where two employees have the same perceived promotabilty, they will have the same probability to get promoted. To make room for promoted agents, a random proportion of the individuals at each level of the hierarchy leave the organization. These spots are then recursively filled until the bottom of the hierarchy is reached. At this point, new agents are then created and "hired" into the entry level of the company. These new agents are equally likely to be men or women.

In this *unbiased model*, there are no differences between men and women: They are equally likely to begin with a given level of promotability, to succeed or fail on projects, to receive stretch projects, and to leave the company. We introduce our model of how the glass ceiling arises through two experiments that extend this unbiased model. First, we introduce six specific, empirically observed ways in which gender discrimination at the interpersonal level manifests in the workplace. Second, we propose a mechanism through which interpersonal gender discrimination is tied to gendered social norms at the macro and meso levels.

Our first modification of the unbiased model introduces six gender biases that have significant empirical support.

First, women's errors and failures on projects are penalized more than men's. For instance, women surgeons experience greater decrease in referrals after a bad outcome: A male surgeon has to have three patient deaths to be penalized the way a female surgeon is penalized after one patient death (Sarsons 2017). We model this gendered penalty by having women agents lose approximately 2 percent more credit than men do for a failed project (see the following).

Second, women's successes on projects are valued less than men's (Bowen, Swim, and Jacobs 2000; Castilla 2008; Eagly, Makhijani, and Klonsky 1992; Moss-Racusin et al. 2012; Swim et al. 1989; Swim and Sanna 1996). For instance,

[1]The title phrase "Insidious Nonetheless" draws from a description of such small, discrete instances by Lenore Blum, who in renouncing her position at Carnegie Mellon University, noted that "Subtle biases and micro-aggressions pile up, few of which on their own rise to the level of 'let's take action,' but are insidious nonetheless. Speak up and you're labeled 'difficult'" (Certo 2018).

in a randomized double-blind study, Moss-Racusin et al. (2012) found that when evaluating candidates for a lab manager position, science faculty at research institutions assigned lower competence values to female applicants than identical male applicants. We model this gendered penalty by having women agents receive approximately 2 percent less credit than men do for a successful project.

Third, women are penalized for exhibiting nonaltruistic behavior (Fanning and David Piercey 2014). Women are seen more unfavorably when they depart from behaviors considered to be stereotypically feminine, such as self-promoting (Rudman 1998). We model this by assigning some percentage of women (here, 10 percent) to complain about receiving less credit than men on project successes. This, in turn, leads them to lose additional credit (here, they receive only 90 percent of the credit they would have gotten had they not complained.

Fourth, women receive fewer opportunities for growth. Women often receive fewer assignments that allow them to develop new skills and report having less access to challenging assignments (King et al. 2012). For example, the American Bar Association found that 44 percent of women of color and 39 percent of white women reported being passed over for desirable assignments in law firms, compared to 2 percent of white men (Rhode 2017). Here, we model this as a requirement that women have 20 percent more successes than a man to be eligible to receive a stretch project.

Fifth, women receive more blame when a mixed-gender team fails (Egan, Matvos, and Seru 2017; Haynes and Lawrence 2012). For instance, participants who receive information about a group's failure assign more blame to women (Haynes and Lawrence 2012). Here, we model this by having women lose approximately 2 percent more credit than their male teammates when a group project fails.

Sixth, women receive less credit in mixed-gender teams (Heilman and Haynes 2005; Sarsons 2017; Sarsons et al. 2021). For example, coauthoring a paper benefits women economists less than it does men: Each coauthored paper increases men's probability of achieving tenure 8.2 percent but increases women's probability of achieving tenure by 5.6 percent (Sarsons et al. 2021). Here, we model this by having women gain approximately 2 percent less credit than their male teammates when a group project succeeds.

These gender biases have empirical support primarily at the interpersonal level. Our model assumes that promotion decisions are made by individuals and that those decisions are a function of gender bias in credit allocation that is accumulated via these six mechanisms. We implement each of these mechanisms into our model as noted briefly previouslu (see the Detailed Methods section at the end of the article for further details).

Most importantly, the six mechanisms vary in their effects. For instance, stretch projects in our model count for 3 times as much as a typical project, so stretch project success can rapidly drive individuals up the corporate hierarchy. In contrast, discounted rewards for women on projects have very small impacts. Thus, a single instance of this form of bias at a single point in time has a minimal effect.

More specifically, to allocate credit for project success and failures, we first assume that the credit c that an agent receives for a project is randomly drawn from a normal distribution. Following Martell et al. (1996), we then vary the percentage of variance in credit received that is explained by gender. We can then use results from prior empirical work to guide the quantity used in our simulation; in particular, we rely, like Martell et al. (2012), on a meta-analysis from Barrett and Morris (1993) that states gender accounts for approximately 1 percent to 5 percent of the variance in hiring decisions. In our model, we fix a parameter r^2, which represents this variance quantity, to .022. This means we assume that gender explains approximately 2 percent of the variation in credit allocation, about half of what Barrett and Morris (1993) found in their study.

Prior empirical work thus guides our parameter settings for how much gender bias impacts credit allocation for project success and failure. For other parameters, however, no empirical evidence we are aware of exists for calibration. Because of this and because a range of values are possible even for empirically informed parameters, we provide results for a range of other parameter settings in the Online Appendix in the Online supplement. While changing model parameters of course changes absolute measures, unless otherwise noted in the following, qualitative findings are consistent across the range of parameter settings we considered.

Results

Figure 1 shows that the interpersonal acts of discrimination we model lead to a glass ceiling effect. In the unbiased model, each level of the corporate hierarchy shows gender parity, with men and women both making up 50 percent of the employees (left-most plot in Figure 1). In contrast, with all of the mechanisms introduced into the model (right-most plot in Figure 1), men dominate upper levels of the corporate hierarchy, leaving a preponderance of women at the lowest levels. Specifically, in the condition where all six bias mechanisms are applied (the right-most subplot of Figure 1), 84 percent of agents at the top of the corporate hierarchy are men. This finding is comparable to the numbers used in the simulation work from Kogut, Colomer, and Belinky (2014), described further in the following, who find that 8.3 percent of directors and 9.1 percent of directors in the "top 500 firms by market cap" are women.

However, not all mechanisms we implement have the same impact; the most significant impacts come from mechanisms that are small but frequently applied. Figure 2 shows

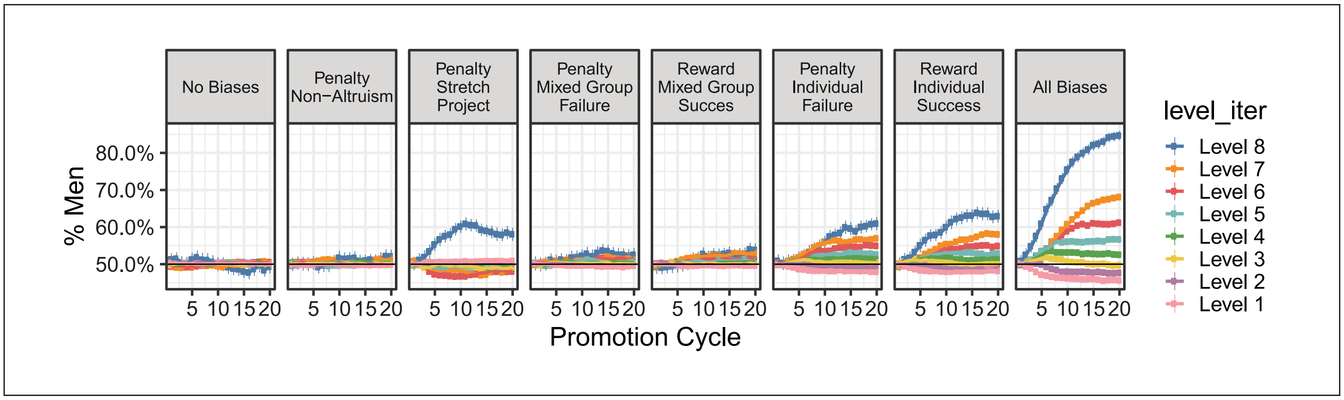

Figure 1. The percentage of men (*y* axis) at each level of the corporate hierarchy (different colors) at each simulated promotion cycle (*x* axis). Different subplots show results for simulations without any empirically validated biases (left-most), with all of these (right-most), or with each individually (middle subplots; results for biases 1 through 6 are shown from left to right). Error bars represent confidence intervals from 300 randomly initialized simulation runs.

that the interpersonal acts of discrimination with the strongest effects on gender disparities were those that had *most frequently been applied*, rather than those with the largest effects on individual agents. As an example of the latter case, differences in growth opportunities via stretch projects—which significantly alter career trajectories, but only for a small number of individuals—impacted gender disparities at the top of the corporate ladder (because successful stretch projects shot individual agents to the top) but were too infrequent in our model to reshape disparities at all levels. Figure 2 also suggests that women who reach high levels are affected more by devaluation for their successes than by penalties for failed projects. This result is explained by the fact that women at higher levels of the hierarchy are more successful (by chance, in our simulation) than women at lower levels. In addition, we find that women at high levels of the corporate hierarchy have a greater track record of successes than their male counterparts (see Online Appendix, Figure A7, in the Online supplement).

Incorporating Social Norms

We have shown that enactment of gender bias at the interpersonal level can result in a glass ceiling for women. However, our model to this point does not express clear assumptions about *why* bias exists in the first place. In this section, we provide such a mechanism based on the existence of gendered social norms at the meso and macro levels. Our starting point is the empirical observation that fewer women in an environment correlates with increased gender discrimination. In management, in the Israeli army, among law students, and in blue-collar work groups, a greater proportion of men results in more bias against women (Lortie-Lussier and Rinfret 2002; Pazy and Oron 2001; Sackett, DuBois, and Noe 1991; Spangler, Gordon, and Pipkin 1978). Prior work has expressed this empirical observation using a mathematical equation that purports that the degree of interpersonal discrimination at one level of the corporate hierarchy changes with the proportion of women at the level above (Robison-Cox et al. 2007). As gender disparities increase, then, gender discrimination does as well, rippling downward throughout the organization.

However, this modeling assumption does little to address claims of "reverse discrimination." That is, such a model must either assume that gender imbalances that *favor women* should result in discrimination *against men* or make the assumption that such reverse discrimination simply cannot exist. The latter claim is unsatisfying theoretically because no underlying mechanism is suggested. However, it is also more consistent with reality. In the few settings where women dominate higher levels of the corporate hierarchy, there is little evidence of men's promotion abilities being impacted. Instead, while women's lack of representation in certain occupations exacerbates disadvantage, men, namely, heterosexual white men, when in short supply, enjoy a *glass escalator*, where they are put on a fast track to advanced positions (Budig 2002; Wingfield 2009), and their evaluation is not affected by their proportion (Pazy and Oron 2001). The preponderance of male school superintendents is one such example (Brunner and Kim 2010).

Our model provides a mechanism that explains both how organizational gender disparities increase gender discrimination and how this can apply only for women. To do so, we draw from scholarship on race and organizations and model how the degree of interpersonal discrimination (and thus resource allocation; Ray 2019) within an organization is a function of social norms that are both internal to (*meso-level norms*) and external to (*macro-level norms*) the organization. We focus here only on project evaluations but note that the model can easily be extended to other interpersonal biases we study as well.

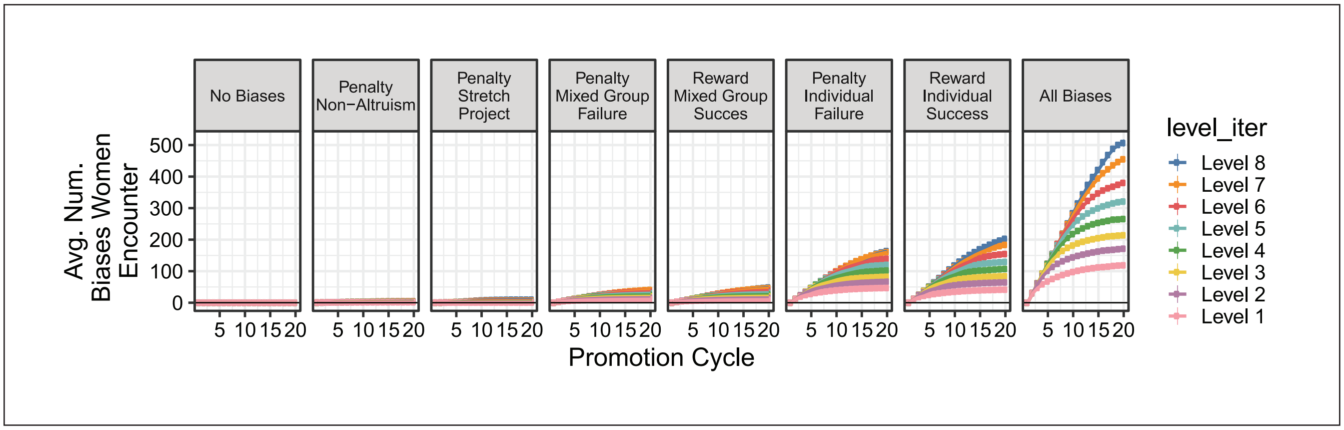

Figure 2. The average number of biases that female agents encounter (*y*-axis) at each level of the corporate hierarchy (different colors) at each simulated promotion cycle (*x*-axis). Different subplots show results for simulations without any empirically validated biases (left-most), with all of these (right-most), or with each individually (middle subplots; results for biases 1 through 6 are shown from left to right). Error bars represent confidence intervals from 300 randomly initialized simulation runs.

More specifically, we introduce the following mathematical model that defines the proportion of variance in project evaluations that is explained by gender:

$$r_i^2 = w \cdot B_{meso,i} + (1-w) \cdot B_{macro} \tag{1}$$

$$B_{meso,i} = \frac{P_{i+1} - 0.5}{P_m - 0.5} \cdot B_{macro}. \tag{2}$$

Here, r_i^2 represents the proportion of variance that gender explains in project credit allocation at level i of the corporate hierarchy. The parameter r_i^2 is a weighted sum of two quantities, where the weight w is also a parameter of the model. The first quantity in the weighted sum is a macro-level norm B_{macro}. This parameter represents an assumption about the variance in project evaluations that would be explained by gender bias if social norms about gender were aligned only with societal biases. The second is a meso-level norm $B_{meso,i}$, which represents the proportion of variance in project evaluations that would be explained by gender if norms were impacted by organizational structure. The value of $B_{meso,i}$ is determined via a formula consisting of P_{i+1}, the proportion of men at level $i+1$ at a given time in the simulation, and P_m, which represents a societal expectation of the percentage of men at a given level of the corporate hierarchy. This value is then multiplied by B_{macro}.

Under this model, the value of r_i^2 is the nweighted average of meso-level and macro-level norms, and the model parameter w encodes the modeler's belief about the relative importance of company-structure-informed social norms compared to societal expectations. Equations 1 and 2 less obviously encode two other core assumptions:

- When the proportion of men in level $i+1$ of the company is .5, $B_{meso,i}$ will be 0: We assume that gender bias driven by social norms within the organization (i.e., meso-level norms) drops to 0 when gender equity is reached.
- When the proportion of men at level $i+1$ of the company is the same as the expected proportion given societal norms, $B_{meso,i}$ will represent the same value as the external norms B_{macro}. We assume that B_{macro} is an accumulation of norms from myriad gendered hierarchies across society. As such, gender norms in levels of a company hierarchy where employee gender distributions match societal expectations (P_m) should mirror the average societal norm, B_{macro}.

We set $B_{macro} = .044$ and initialize all levels of the simulated corporation to have 80 percent women to represent a women-dominated organization. We then explore how glass escalators do or do not emerge under different assumptions about the values of P_m (expected proportion of men at each level of the corporate hierarchy) and w (relative weight of meso vs. macro norms). Figure 3 shows that for fixed values of P_m (within each row of the figure), when we assume that macro-level norms have less influence (from left to right), men tend to face more interpersonal discrimination in women-majority companies. Because these reverse biases are rarely observed empirically, we argue that a model that considers social norms only at the meso level is incomplete. Instead, Figure 3 shows that only models that incorporate both meso and macro norms, and more specifically, models that heavily weight societal-level norms relative to norms attributable to gender disparities within organizations, display evidence of the empirically observed glass escalator effect. Notably, however, this effect changes as

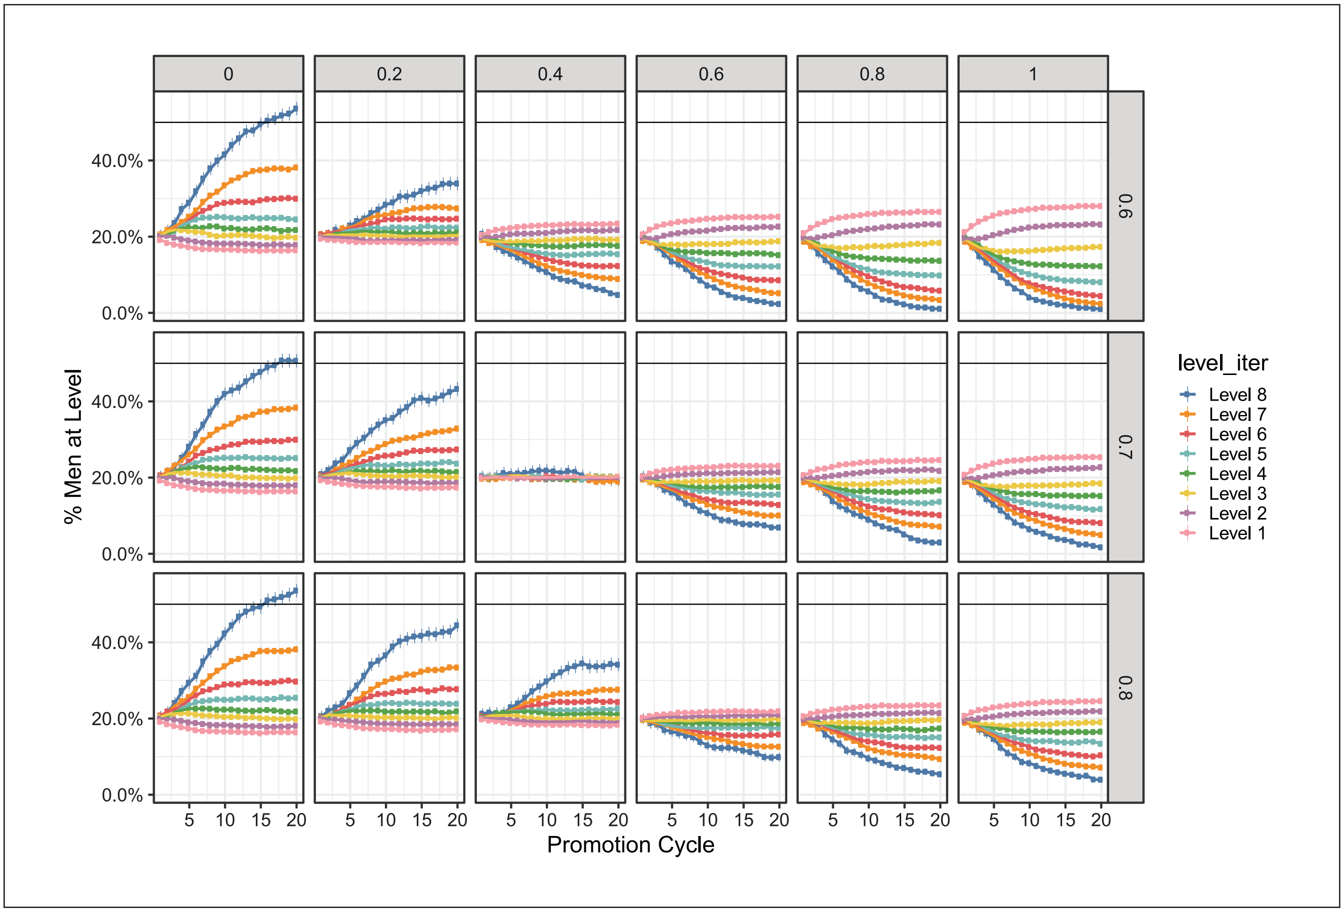

Figure 3. The percentage of employees that are men (y-axis) at each level of the corporate hierarchy (different colors) at each simulated promotion cycle (x-axis). Different columns show results for simulations where we vary the strength of meso-level norms relative to macro-level norms (i.e., the model parameter w). A value of 0 represents a model in which only macro (societal) norms influence agent decisions and 1 represents that only meso (organizational) norms impact agent decisions. Different rows show results for simulation where we vary societal expectations of the proportion of men at each level (i.e., P_m). All simulations here assume that at the onset of the simulation, all levels of the corporate hierarchy are made up of 80 percent women (i.e., that $P_{male} = .2$). Other parameters used are introduced in the Online Appendix, Table A1 and Online Appendix, Table A4 in the Online supplement. Error bars represent confidence intervals from 300 randomly initialized simulation runs. The black horizontal bar represents 50 percent men as a reference point.

we vary P_m (i.e., across the rows of Figure 3). As P_m increases (from top to bottom), meso-level norms must have a larger assumed effect for men to face reverse discrimination.

Ultimately, then, we show that reverse discrimination can occur, but under very specific conditions. Put another way, r_i^2 can be negative, in which case we model reverse discrimination as in prior work. However, the parameters w and P_m mitigate this possibility—if meso-level norms matter little in comparison to macro-level norms (i.e., if w is small) or societal expects a predominance of men at a particular level of the corporate hierarchy (P_m is close to 1), reverse discrimination is unlikely.

Implications for Intervention

A common approach to mitigating gender disparities in organizations is to implement a quota-based system that enforces rules about promotions based on gender (Pande and Ford 2012). Here, we simulate the effects of a quota-based intervention using our model. After 7 promotion cycles without intervention, a quota system is introduced to our simulated company for six promotion cycles (our finding is robust to the number of promotion cycles, see Online Appendix, Figure A8 and Online Appendix, Figure A9 in the Online supplement). The quota intervention we assume is one where rules on promotions are enforced that target a goal of having K percent of each level of the company above the entry level be women. We vary the value of K to understand how different degrees of intensity of quota-based intervention would impact the gender structure of the corporation in the long term under existence of macro-level norms. We then vary the assumed strength of meso-level norms, relative to macro norms, within the company.

We find that gender disparities in our simulated organizations will return over time if company gender norms are at all displaced by gender-biased macro norms, even with quota levels as high as 70 percent. These results are shown in Figure 4;

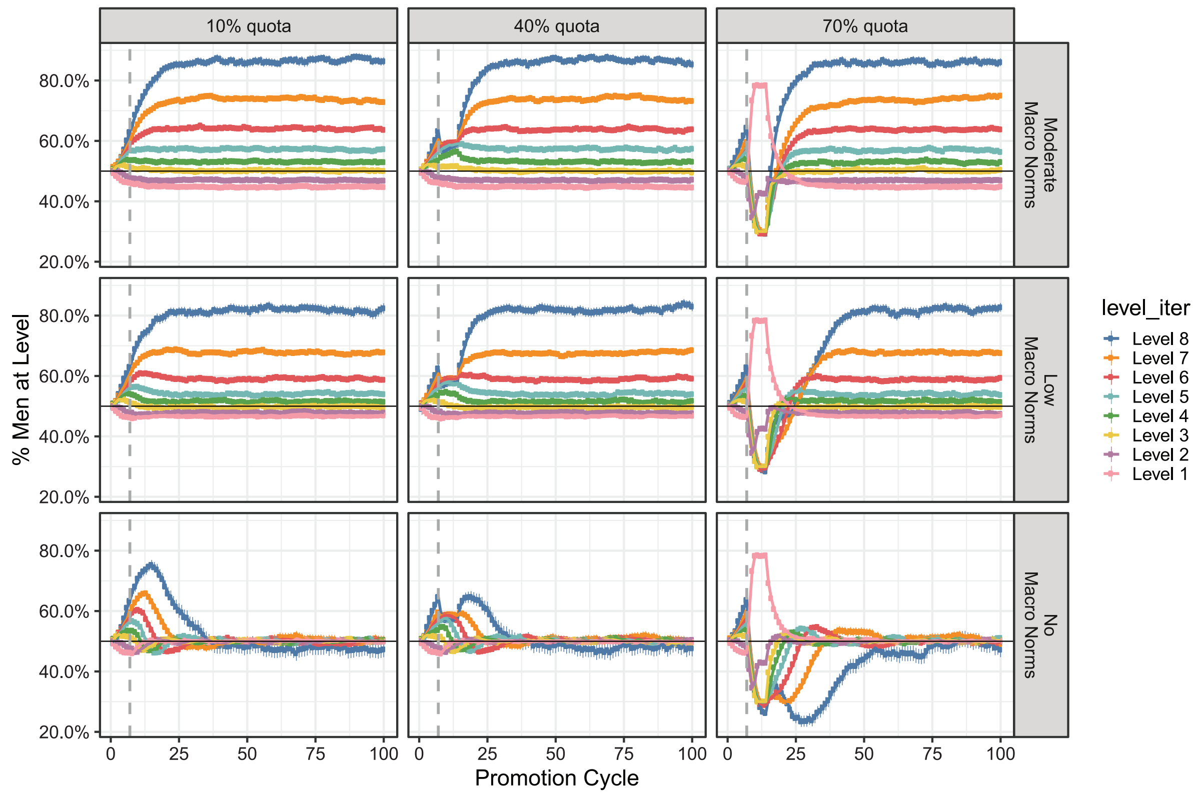

Figure 4. The percentage of employees that are men (*y*-axis) at each level of the corporate hierarchy (different colors) at each simulated promotion cycle (*x*-axis). Different rows show results for simulations where we vary the strength of meso-level norms relative to macro-level norms (i.e., the model parameter w). Values are .4, .7, and 1 for Moderate, Low, and No Macro Norms, respectively. A value of 0 represents a model in which only macro (societal) norms influence agent decisions, and 1 represents that only meso (organizational) norms impact agent decisions. Different columns show results for different degrees of intensity of quota-based intervention (i.e., the model parameter K). All simulations here assume that quota-based intervention is carried out for six Promotion Cycles (i.e., $I_{range} = [168, 312]$). Error bars are confidence intervals from 300 randomly initialized simulation runs.

only in the last row of figures, where macro-level norms have no impact on promotion decisions, do we observe gender disparities gradually decrease after reaching the peak. At face value, these findings conflict with those from the agent-based model of Kogut et al. (2014), who find that quotas as small as 20 percent can induce forms of structural equality across genders. However, Kogut et al. (2014) note that their model, while providing important insights into the gendered social network structure of board directors, does not account for the role of societal norms and/or beliefs about gender. Our work thus provides a caution to their findings, noting that one must, as previously mentioned, make generous assumptions about the malleability of macro-level gender norms for quota-based interventions to sustain long-term impacts on the gendered nature of corporate structures.

Discussion

Gender disparities in organizations emerge from a complex, dynamic social system (Martell et al. 2012). Prior agent-based models have shown how a variety of mechanisms, such as career interruptions and variable attrition rates (Bullinaria 2018; Momennejad et al. 2019; Robison-Cox et al. 2007), can create gender disparities in these complex organizational systems. Most recently, Momennejad et al. (2019) simulate the costs to individuals and institutions of sexist comments and objections to those comments in meetings, finding interrelationships between structural and learning effects. Their work shows how social learning can prevent structural interventions from being effective.

The present work extends these efforts. We provide a concrete mechanism through which empirical observations of gender discrimination at the interpersonal level can be embedded into a model of complex organizational systems. Doing so paves a path toward better integration of empirical and simulation results in the study of the glass ceiling. Furthermore, prior work has largely focused on identifying the effect sizes of mechanisms for gender bias. Our work instead models both effect size and the *frequency* with which these small interpersonal acts of gender bias are enacted.

This is important because we find in our simulations that biased evaluations that produce small effects but that occur at frequent intervals over a period of time may be the most consequential in producing gender disparities. Finally, we introduce a new hierarchical model of how organizational and societal norms combined to create gender discrimination. In doing so, we argue via simulation that interventions aimed at reducing gender disparity in organizations must attend to the strength of societal gender norms and the stubbornness of outside influences when devising plans to disrupt gender homogeneity in corporate hierarchies. Critically, then, our model encourages further attention toward radical, societal-level change, or at least changes at the meso level that can be expected to diffuse out to macro-level structures, such as quotas in roles with direct policy implications (Beaman et al. 2009) or quotas in combination with efforts to shift widespread societal norms via, for example, coordination with widespread policy change (Ofosu et al. 2019).

In considering these advancements of our work over prior efforts, however, it is important to also note our limitations. First, while we focus on gender as a binary, we emphasize that gender itself is a continuous and socially constructed system (Ridgeway 2011). Second, while our model could be easily extended to focus on intersectional dimensions of inequality and discrimination, the focus in the present work is on gender and thus does not account explicitly for the intersectional nature of inequality or the ways in which stereotypes associated with other groups interact with gender stereotypes to amplify or dilute biases (Hall et al. 2019). Third, it is difficult to know the true impact of small, continuously applied interpersonal biases. In any case, actual effect sizes will vary by organization and by individuals within organizations. Our model, informed by empirical results, assumes very small effect sizes; in the real world, these may be larger, smaller, or inconsistently applied. While we have provided robustness tests for our modeling assumptions in the Online Appendix in the Online supplement, these are of course subject to similar concerns. Fourth, many factors contribute to any individual's career trajectory beyond those listed here: choices and preferences, workplace family policies, and more. Our model demonstrates only that disparities on the order of magnitude of those seen in the real world can be achieved via the interpersonal mechanisms presented, with full account of the norms on which these interpersonal actions are based.

Despite these limitations, our work serves broader theoretical and policy-oriented goals. With respect to theory, our model provides a link between status construction theory (Mark, Smith-Lovin, and Ridgeway 2009), which focuses on the link between norms and behavior, and Ray's (2019) theory of racial inequality emphasizing how culture, resources, and ideology interact at the micro, meso, and macro levels of analysis. With respect to policy, laws are designed to address either large events that happen infrequently and can be easily attributed to a single actor—for example, overt sexual harassment by a manager—or "pattern and practice" in an organization, for instance, explicitly discriminatory policies. Our model shows, by contrast, how large organizational disparities can occur via that gradual and diffuse impact of many small, even unintentional events, decisions, and evaluations happening frequently over a long period of time. This raises important questions about the location of accountability within organizations and organizational culture and about what role the legal system or even workplace policies can or should play in cases where the biased evaluations are of the sort we model here.

Detailed Methods

We provide here more complete details on the simulation model used in this article. Full parameter details are in the Online Appendix, Tables A1 through A4 in the Online supplement, and full replication materials are available at https://github.com/yuhaodu/workplace_gender_bias.

Further Model Details: Agents

Agent states in our model are constituted by variables that keep track of the number of successful and failed projects this agent has completed and the agent's perceived promotability. Each agent also has a binary attribute for their perceived binarized gender—man or woman.

Further Model Details: Company

We model the same eight-level organization as Martell et al. (1996). Level 8 represents the highest level of the company (i.e., the C-suite executives), and Level 1 represents the lowest level. At the beginning of the simulation, all positions at all levels are seeded with agents who are randomly assigned a gender. As in Martell et al. (1996), the eight levels have 10, 40, 75, 100, 150, 200, 350, and 500 agents, respectively.

The company evolves through a series of project turns. Each project turn can be either a traditional project turn or a stretch project turn. Stretch project turns occur once every 12 turns. On a traditional project turn, 50 percent of agents are randomly assigned to individual projects, the others to group projects. On a stretch project turn, stretch projects will first be assigned to $P_{stretch}$ percent of agents. Then $(1-P_{stretch})*P_{individual}$ percentage of agents receive individual projects, while the rest will be assigned to group projects. In this work, $P_{stretch}=.1$ and $P_{individual}=.5$ for all runs.

After $n_{promotion}$ project turns, the company will carry out one *promotion cycle turn* (here, $n_{promotion}=24$). Promotion cycle turns happen in a sequence of two steps. First, a random P_{leave} percentage of agents at each level of the company leave (here, P_{leave} = 15 percent). Second, the company carries out a series of promotions, where empty positions caused by agents leaving the company are filled by agents who occupy the lower level positions. Agents that are promoted are those that have the highest perceived promotability. Empty positions at lowest level are filled by new agents.

Further Model Details: Projects

There are three kinds of projects in our simulation—individual projects, stretch projects, and group projects. Individual projects and stretch projects are both assigned to a single agent. Group projects are assigned to two agents. All projects have an attribute, c, that is used to determine the amount of credit (blame) given to agents assigned to the project when it succeeds (fails). The value of c is drawn from a normal distribution with mean μ_r and standard deviation σ_r for individual and group projects and from a normal distribution with mean μ_{st} and standard deviation σ_{st} for stretch projects. Simulations in this article are run with $\mu_r = 10, \sigma_r = 1, \mu_{st} = 30, \sigma_{st} = 1$, reflecting an assumption of stretch projects being roughly 3 times as important as the typical project. In our simulation, we make the simplifying assumption that all projects are equally likely to succeed or fail. With no gender bias, if a project succeeds, the perceived promotability of the agents assigned to the project will increase by c; if it fails, the perceived promotability of the agents assigned will decrease by c.

Modeling That Women's Successes (Failures) on Independent and Group Projects Are Valued (Penalized) Less (More) Than Men's

We operationalize devalued success for women on projects using the percentage of variance in project credit that is explained by agent gender. More specifically, model parameters r_i^2, introduced in Equation 1, can be interpreted as the percentage of variance explained by agent gender in a linear regression where the dependent variable is c, the credit the agent (at level i of the company) receives for completing a successful project. In Figure 1, credit received is independent of the agent's level of the company, and thus we discuss a parameter r^2, where $r_i^2 = r^2 \forall i$. For Figure 1, $r^2 = .022$. Practically, this is implemented by setting w to 0 in Equation 1 in the main text and fixing $B_{macro} = .022$.

To explain how gender bias in project credit allocation is implemented, we focus on this level-independent value r^2. The details stated here go through analogously with parameters r_i^2. Implementing percentage variance explained in the simulation requires a variable transformation from r^2 to a raw value, d, that differentiates credit given to women and credit given to men. To do so, we first expand notation, assuming the perceived promotability of a male agent will increase by c upon the completion of a successful individual project, while the perceived promotability of a female agent will increase by only $c - d$. We then derive the appropriate value of d such that this process will result in a particular value of r^2. To do so, note again that the credit of a project is drawn from a normal distribution with mean μ_r and σ_r. Now, define $d = \frac{2 \cdot r}{\sqrt{1 - r^2}}$ such that d represents the standardized mean difference between credit allocated to men and women. Let us now define μ_g and σ_g to represent the mean and standard deviation of project credit allocated to agents with gender g. Via simple derivation, it can be said that $\mu_{male} - \mu_{female} = d \cdot \sqrt{2 \cdot (\sigma^2_{male} + \sigma^2_{female})}$. In turn, $\mu_{male} - \mu_{female} = d$ if we set the σ_r to 1.

Thus, by fixing $\sigma_r = 1$, as we do in the simulation, we can model the fact that gender explains r^2 percent of the variance in credit allocation via the following procedure. First, for a successful project, we sample credit c for this project. The perceived promotability of a male employee will then increase by c, and the perceived promotability of a female employee will only increase by $c - d$. In this way, we can simulate an environment where gender bias accounts for r^2 proportion of the variance.

Note also that the quantity $\frac{d}{c}$ can be understood as the average amount that a man's perceived promotability will increase over and above a woman's for the same successful project. That is, given fixed values for r^2, μ_r, and σ_r, one can compare the raw percentage increase that a male versus a female agent receives in perceived promotability for each successful project completed. Because of this dependence on some unknowable "absolute increase in promotability per project success," the quantity of interest for both our work and Martell et al. (1996) is thus not $\frac{d}{c}$ but r^2. Finally, we note again that it is possible for r_i^2 to be negative. In this case, our simulation code instead models $\mu_{female} - \mu_{male} = d$, effectively encoding so-called reverse discrimination.

Similarly, for failed individual projects, men's perceived promotability decreases by c for failed projects and women's by $c + d$. To model biased allocation of credit in mixed-gender teams for success and failure, we adopt the same procedure as we do for individual projects. The only difference is that we use a different parameter, $\tilde{r}^2$, and analogously $\tilde{r}_i^2$ when level-specific biases are considered.

Modeling That Women Are Penalized for Exhibiting Non-altruistic Behavior

In our model, we assign a percentage of women, P_{com} to occasionally self-promote by complaining about unfairness when they experience bias. Doing so leads to a decrease in their credit score when they engage in this behavior by multiplying a discount factor f_{dis} to their credit. If a female agent engages in self-promotion activity, their credit will change from c to $f_{dis} \cdot c$, where f_{dis} ″ 1. In the simulations presented in the main text, $P_{com} = .1$ and $f_{dis} = .9$.

Modeling That Women Receive Fewer Opportunities for Growth

In our model, at fixed intervals (every $n_{stretch} = 12$ turns in the models in the present work), we introduce stretch projects that provide outsized boosts in perceived promotability. Women need to achieve P_{female} more successful projects than those of the average of qualified men to be assigned stretch projects. In the results presented here, $P_{female} = 20\ percent$; women thus need 20 percent more successes to be considered for stretch projects (we conduct robustness tests on selection of P_{female}. See Appendix, Figure A6). On each stretch project turn, we first rank the agents according to their perceived promotability. The top $P_{stretch}$ percentage of agents are then considered to be *prequalified* for stretch projects. In the results presented here, $P_{strech} = 10\ percent$. From these prequalified agents, we calculate the average number n_{avg} of successful projects that male agents have already finished. Female agents then must have had to finish $n_{avg} \cdot (1 + P_{female})$ successful projects to be qualified for stretch projects.

Modeling the Quota Intervention

The quota-based intervention study we introduce has a single parameter, K, that specifies a quota for the percentage of female agents expected at each level of the company. Thus, if Level $i+1$ has n positions and n_f is the number of female employees at Level i, we will try to promote $n \cdot K\ percent$ – n_f female employees from level i to guarantee that there are at least K percent female employees at Level $i+1$. Other positions at Level $i+1$ are filled by employees who have highest perceived promotability from Level i. We vary K in 10, 40, and 70 to present mild, intermediate, and aggressive quota interventions, respectively.

We evaluate this intervention by further varying two additional parameters. Figure 4, and in the Online Appendix, Figure A8 and in the Online Appendix, Figure A9 in the Online supplement show the results about different ranges of project turns, I_{range}, which determine the project turn on which the intervention starts and the project turn on which the intervention ends. Values of I_{range} [168, 240], [168, 312], and [168, 384] correspond to the three, six, and nine promotion cycles. We also vary the weight of meso-level norms. In all cases, the weight of meso-level norms starts with $w_0 = .5$. Then, at the beginning of the intervention, and through the rest of the simulation, the weight will be altered to w. We set w to .4, .7, and 1, aligning with the "Moderate Macro Norms," "Low Macro Norms," and "No Macro Norms" labels, respectively, of the plot rows in three aforementioned figures.

ORCID iD

Kenneth Joseph https://orcid.org/0000-0003-2233-3976

Supplemental Material

Supplemental material for this article is available online.

References

Babcock, L., and S. Laschever. 2009. Women don't ask. In *Women Don't Ask*. Princeton University Press.

Barrett, G. V., and S. B. Morris. 1993. "The American Psychological Association's Amicus Curiae Brief in Price Waterhouse v. Hopkins." *Law and Human Behavior* 17(2):201–15.

Beaman, L., R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova. 2009. "Powerful Women: Does Exposure Reduce Bias?" *The Quarterly Journal of Economics* 124(4):1497–540.

Bear, J. B. 2021. "Forget the 'Mommy Track': Temporal Flexibility Increases Promotion Aspirations for Women and Reduces Gender Gaps." *Psychology of Women Quarterly* 45(3):294–307. doi:10.1177/03616843211003070.

Berdahl, J. L., and C. Moore. 2006. "Workplace Harassment: Double Jeopardy for Minority Women." *Journal of Applied Psychology* 91(2):426–36.

Bertrand, M. 2018. "Coase Lecture – The Glass Ceiling." *Economica* 85(338):205–31. doi:10.1111/ecca.12264.

Bowen, C. C., J. K. Swim, and R. R. Jacobs. 2000. "Evaluating Gender Biases on Actual Job Performance of Real People: A Meta-analysis." *Journal of Applied Social Psychology* 30(10):2194–215.

Brunner, C. C., and Y. L. Kim. 2010. "Are Women Prepared to Be School Superintendents? An Essay on the Myths and Misunderstandings." *Journal of Research on Leadership Education* 5(8):276–309.

Budig, M. J. 2002. "Male Advantage and the Gender Composition of Jobs: Who Rides the Glass Escalator?" *Social Problems* 49(2):258–77.

Bullinaria, J. 2018. "Agent-Based Models of Gender Inequalities in Career Progression." *Journal of Artificial Societies and Social Simulation* 21(3):7. doi:10.18564/jasss.3738.

Carley, K. 1991. "A Theory of Group Stability." *American Sociological Review* 56(3):331–54. doi:10.2307/2096108.

Castilla, E. J. 2008. "Gender, Race, and Meritocracy in Organizational Careers". *AJS: American Journal of Sociology* 113(6):1479–526. doi:10.1086/588738.

Certo, T. 2018. "Lenore Blum Shocked the Community with Her Sudden Resignation from CMU. Here She Tells Us Why." *Next Pittsburgh*, September 6. https://nextpittsburgh.com/features/lenore-blum-speaks-out-about-sexism-in-the-workplace/.

Collischon, M. 2021. "Personality Traits as a Partial Explanation for Gender Wage Gaps and Glass Ceilings." *Research in Social Stratification and Mobility* 73:100596. doi:10.1016/j.rssm.2021.100596.

Cotter, D. A., J. M. Hermsen, S. Ovadia, and R. Vanneman. 2001. "The Glass Ceiling Effect." *Social Forces* 80(2):655–81.

Eagly, A. H., M. G. Makhijani, and B. G. Klonsky. 1992. "Gender and the Evaluation of Leaders: A Meta-analysis." *Psychological Bulletin* 111(1):3–22.

Egan, M. L., G. Matvos, and A. Seru. 2017. "When Harry Fired Sally: The Double Standard in Punishing Misconduct." Technical report, National Bureau of Economic Research, Cambridge, MA.

Fanning, K., and M. David Piercey. 2014. "Internal Auditors' Use of Interpersonal Likability, Arguments, and Accounting Information in a Corporate Governance Setting." *Accounting, Organizations and Society* 39(8):575–89. doi:10.1016/j.aos.2014.07.002.

Filer, R. 1983. "Sexual Differences in Earnings: The Role of Individual Personalities and Tastes." *Journal of Human Resources* 18(1):82–99.

Gilbert, N. 2007. *Agent-Based Models*. Thousand Oaks, CA: Sage Publications Inc. ISBN 978-1-4129-4964-4.

Goldin, C., and L. F. Katz. 2016. "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation." *Journal of Labor Economics* 34(3):705–46. doi:10.1086/685505.

Hall, E. V., A. V. Hall, A. D. Galinsky, and K. W. Phillips. 2019. "MOSAIC: A Model of Stereotyping through Associated and Intersectional Categories." *Academy of Management Review* 44(3):643–72.

Harrison, J. R., Z. Lin, G. R. Carroll, and K. M. Carley. 2007. "Simulation Modeling in Organizational and Management Research." *Academy of Management Review* 32(4):1229–45.

Haynes, M. C., and J. S. Lawrence. 2012. "Who's to Blame? Attributions of Blame in Unsuccessful Mixed-Sex Work Teams." *Basic and Applied Social Psychology* 34(6):558–64. doi:10.1080/01973533.2012.727312.

Heilman, M. E., and M. C. Haynes. 2005. "No credit Where Credit Is Due: Attributional Rationalization of Women's Success in Male-Female Teams." *The Journal of Applied Psychology* 90(5):905–16. doi:10.1037/0021-9010.90.5.905.

King, E. B., W. Botsford, M. R. Hebl, S. Kazama, J. F. Dawson, and A. Perkins. 2012. "Benevolent Sexism at Work: Gender Differences in the Distribution of Challenging Developmental Experiences." *Journal of Management* 38(6):1835–66. doi:10.1177/0149206310365902.

Kogut, B, J. Colomer, and M. Belinky. 2014. "Structural Equality at the Top of the Corporation: Mandated Quotas for Women Directors." *Strategic Management Journal* 35(6):891–902.

Lortie-Lussier, M., and N. Rinfret. 2002. "The Proportion of Women Managers: Where Is the Critical Mass?" *Journal of Applied Social Psychology* 32(9):1974–91. doi:10.1111/j.1559-1816.2002.tb00268.x.

Mark, N. P., L. Smith-Lovin, and C. L. Ridgeway. 2009. "Why Do Nominal Characteristics Acquire Status Value? A Minimal Explanation for Status Construction." *American Journal of Sociology* 115(3):832–62.

Martell, R. F., C. G. Emrich, and J. Robison-Cox. 2012. "From Bias to Exclusion: A Multilevel Emergent Theory of Gender Segregation in Organizations." *Research in Organizational Behavior* 32:137–62.

Martell, R. F., D. M. Lane, and C. Emrich. 1996. "Male-Female Differences: A Computer Simulation." *American Psychologist* 51(2):157–58. doi:10.1037/0003-066X.51.2.157.

Miller, C. C., K. Quealy, and M. Sanger-Katz. 2018. "The Top Jobs Where Women Are Outnumbered by Men Named John." *The New York Times*, April 24. https://www.nytimes.com/interactive/2018/04/24/upshot/women-and-men-named-john.html.

Momennejad, I., S. Sinclair, and M. Cikara. 2019. "Computational Justice: Simulating Structural Bias and Interventions." bioRxiv. doi:10.1101/776211.

Moss-Racusin, C. A., J. F. Dovidio, V. L. Brescoll, M. J. Graham, and J. Handelsman. 2012. "Science Faculty's Subtle Gender Biases Favor Male Students." *Proceedings of the National Academy of Sciences* 109(41):16474–79.

Ofosu, E. K., M. K. Chambers, J. M. Chen, and E. Hehman. 2019. "Same-Sex Marriage Legalization Associated with Reduced Implicit and Explicit Antigay Bias." *Proceedings of the National Academy of Sciences* 116(18):8846–51.

Pande, R., and D. Ford. 2012. *Gender Quotas and Female Leadership*. Washington, DC: World Bank.

Pazy, A., and I. Oron. 2001. "Sex Proportion and Performance Evaluation among High-Ranking Military Officers." *Journal of Organizational Behavior* 22(6):689–702. doi:10.1002/job.109.

Pettit, B., and J. Hook. 2005. "The Structure of Women's Employment in Comparative Perspective." *Social Forces* 84(2):779–801. doi:10.1353/sof.2006.0029.

Ray, V. 2019. "A Theory of Racialized Organizations." *American Sociological Review* 84(1):26–53.

Reuben, E., P. Sapienza, and L. Zingales. 2015. "Taste for Competition and the Gender Gap among Young Business Professionals." Technical report, National Bureau of Economic Research, Cambridge, MA.

Rhode, D. L. 2017. *Women and Leadership*. Oxford, UK: Oxford University Press.

Ridgeway, C. L. 2011. *Framed by Gender: How Gender Inequality Persists in the Modern World*. Oxford, UK: Oxford University Press.

Robison-Cox, J. F., R. F. Martell, and C. G. Emrich. 2007. "Simulating Gender Stratification." *Journal of Artificial Societies and Social Simulation* 10(3):1–8.

Rudman, L. A. 1998. "Self-Promotion as a Risk Factor for Women: The Costs and Benefits of Counterstereotypical Impression Management." *Journal of Personality and Social Psychology* 74(3):629–45.

Sackett, P. R., C. L. DuBois, and A. W. Noe. 1991. "Tokenism in Performance Evaluation: The Effects of Work Group Representation on Male-Female and White-Black Differences in Performance Ratings." *Journal of Applied Psychology* 76(2):263–67. doi:10.1037/0021-9010.76.2.263.

Sarsons, H. 2017. "Interpreting Signals in the Labor Market: Evidence from Medical Referrals." Job Market Paper. https://scholar.harvard.edu/files/sarsons/files/sarsons_jmp.pdf.

Sarsons, H., K. Gërxhani, E. Reuben, and A. Schram. 2021. "Gender Differences in Recognition for Group Work." *Journal of Political Economy* 129(1):101–47. doi:10.1086/711401.

Schubert, R., M. Gysler, M. Brown, and H. W. Brachinger. 2000. "Gender Specific Attitudes towards Risk and Ambiguity: An Experimental Investigation." Technical report, Economics Working Paper Series, No. 00/17, ETH Zurich, CER-ETH - Center of Economic Research, Zurich.

Semykina, A., and S. J. Linz. 2007. "Gender Differences in Personality and Earnings: Evidence from Russia." *Journal of Economic Psychology* 28(3):387–410. doi:10.1016/j.joep.2006.05.004.

Spangler, E., M. A. Gordon, and R. M. Pipkin. 1978. "Token Women: An Empirical Test of Kanter's Hypothesis." *American Journal of Sociology* 84(1):160–70. doi:10.1086/226745.

Stockdale, M. S., and G. Bhattacharya. 2009. "Sexual Harassment and the Glass Ceiling." Pp. 171–99 in *The Glass Ceiling in the 21st Century: Understanding Barriers to Gender Equality*, edited by M. Barreto, M. K. Ryan, and M. T. Schmitt. Washington, DC: American Psychological Association. doi:10.1037/11863-008.

Swim, J., E. Borgida, G. Maruyama, and D. G. Myers. 1989. "Joan McKay Versus John McKay: Do Gender Stereotypes Bias Evaluations?" *Psychological Bulletin* 105(3):409–29.

Swim, J. K., and L. J. Sanna. 1996. "He's Skilled, She's Lucky: A Meta-analysis of Observers' Attributions for Women's and Men's Successes and Failures." *Personality and Social Psychology Bulletin* 22(5):507–19.

Williams, J., and N. Segal. 2003. "Beyond the Maternal Wall: Relief for Family Caregivers Who Are Discriminated Against on the Job." *Harvard Women's Law Journal* 26:77.

Wingfield, A. H. 2009. "Racializing the Glass Escalator: Reconsidering Men's Experiences with Women's Work." *Gender & Society* 23(1):5–26. doi:10.1177/0891243208323054.

Author Biographies

Yuhao Du is a doctoral student in the Department of Computer Science and Engineering at the University at Buffalo. He is a computational social scientist whose research has been published in the *International Conference on the Web and Social Media (ICWSM)* and has been featured in the *New York Times*. His research interests include simulation of sociotechnical systems, and the analysis of fairness and accountability in machine learning algorithms.

Jessica Nordell is a writer and science journalist with a background in physics and poetry. Her first book *THE END OF BIAS: A BEGINNING* (Metropolitan Books/Holt Macmillan) is a solutions-based odyssey through the science and psychology of unexamined bias and discrimination and a story-driven tour through evidence-based approaches to ending it. Her writing has appeared in the *New York Times*, the *Atlantic, Slate*, the *New Republic*, and the *Washington Post*. She received her B.A. in physics from Harvard and M.F.A. in poetry from the University of Wisconsin-Madison.

Kenneth Joseph is an Assistant Professor in the Department of Computer Science and Engineering at the University at Buffalo. He received his Ph.D. from the Societal Computing program at Carnegie Mellon University, was a postdoctoral researcher at the Network Science Institute at Northeastern University, and a former fellow at Harvard's Institute for Quantitative Social Science. His research has appeared in a variety of outlets, including *The Journal of Mathematical Sociology*, a number of top computer science conferences, and *Science*. His work has also been covered in popular outlets including *The Washington Post* and *The Chicago Tribune*.