

**Diversity, Polarization, and Recruiting**

by

William Hurst

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
(Business Administration and Political Science)  
in the University of Michigan  
2024

Doctoral Committee:

Professor Jordan Siegel, Co-Chair  
Professor Robert Franzese, Co-Chair  
Assistant Professor Justin Frake  
Dr. Mara Ostfeld  
Associate Professor Christopher Rider  
Professor Tanya Rosenblat

William Hurst

[hursre@umich.edu](mailto:hursre@umich.edu)

ORCID iD: [0000-0002-2589-2805](https://orcid.org/0000-0002-2589-2805)

© William R. Hurst 2024

## **Acknowledgements**

Faculty, friends, and family have been instrumental to this dissertation. I have benefited immensely from the orientation of my Strategy dissertation members, including Justin Frake, Chris Rider, and Jordan Siegel. Felipe Csaszar, Cheng Gao, Derek Harmon, Michael Jensen, Jim Westphal, Brian Wu, and Maggie Zhou have also provided much-needed guidance along the way. A unique part of my Michigan experience has been the opportunity to pursue a joint PhD in political science. I have benefited from the insight of political scientists Rob Franzese and Mara Ostfeld, who kindly agreed to serve on my committee, as well as Allen Hicken, who originally floated the idea of a dual degree. I also am grateful to Tanya Rosenblat who gave much needed encouragement early on and has been a helpful source of feedback ever since.

I have also benefited greatly from my PhD classmates. Mike Ricks patiently guided me through tough econometrics training, has provided exceptional feedback on drafts, and is now a close friend and co-author. Diana Jue-Rajasingh, with whom I braved the job market, was the first Michigan Strategy PhD student I met and has also become a friend, co-author, and great source of feedback. Since my first year, Ronnie Lee has been a constant source of encouragement, support, and advice. He inspires me with his work ethic. Shonita Black has become a close friend, in no small part from a bond formed during a long, unexpected drive to and from Indianapolis (thank you!). I also benefited from the friendship and support of other PhD students including Yun Ha

Cho, Christine Choi, Susie Choi, David Chung, Mana Heshmati, Harsh Ketkar, Raji Kunapuli, Derek Leif, Cha Li, Kyle McCullers, Brenda Myung, Aseem Sinha, and Weikun Yang.

Finally, and most importantly, I benefited from the support of family. My parents, Grant and Jeanette, have always been a phone call away, ready to listen to whatever victory or disappointment I was ready to share. My siblings, and their extremely cute children, have been a source of joy and provided much needed reminders when I began to over-prioritize professional life. Since the start of the PhD, Winston, Phil, and I have had (at least) annual backcountry fishing trips into the Wind River, Sierra Nevada, and/or Teton mountains. Memories from these trips are sacred. No one deserves more acknowledgement than my wife Kallan. We met within my first few months at Michigan. She is not only my emotional bedrock but has been a sounding board for research ideas as she has embarked on her own PhD. Michigan will always hold a special place in our hearts. Our Australian shepherd, Laska, has been a beloved, constant companion (literally laying at my feet) for most of the last three years.

## Table of Contents

Acknowledgements .....	ii
List of Tables .....	viii
List of Figures .....	xi
Abstract .....	xiv
Chapter 1 Brief Introduction and Overview .....	1
1.1 Chapter 1: Countervailing Claims Pro-Diversity Responses to Stigma by Association Following the Unite the Right Rally .....	3
1.2 Chapter 2: Workplace Backlash? Workforce Diversity, Status Threat, and the Contractionary Effects of Pro-Diversity Claims .....	5
1.3 Chapter 3: The Effect of Flatter Hierarchy on Applicant Pool Gender Diversity: Evidence from Experiments .....	7
1.4 Looking Forward: Partisan De-diversification in the workplace .....	9
1.5 Works Cited .....	11
Chapter 2 Countervailing Claims Pro-Diversity Responses to Stigma by Association Following the Unite the Right Rally .....	17
2.1 Introduction .....	18
2.2 Theory .....	20
2.2.1 Stakeholder Responses to Sociopolitical Positioning .....	21
2.2.2 Combating Sociopolitical Stigma with Countervailing Claims .....	24
2.3 Empirical Context: Unite the Right Rally in Charlottesville, Virginia .....	27
2.3.1 Sociopolitical Stigma for Charlottesville Employers .....	28
2.3.2 Countervailing Pro-Diversity Claims .....	30
2.4 Econometric Analysis .....	34
2.4.1 Countervailing Pro-Diversity Claims .....	34
2.4.2 Appropriate Level of Analysis .....	37
2.4.3 Summary Statistics .....	39
2.4.4 Difference-in-Differences Analysis .....	40
2.4.5 Evaluating Plausibility of the Parallel Trends Assumption .....	41
Parallel trends sensitivity analyses .....	42
2.4.6 Considering Alternative Control Geographies .....	42
2.4.7 Employer-Level Analysis .....	45
2.4.8 Stable Unite Treatment Value Assumption Violations .....	46

2.5 Examining the Proposed Mechanism.....	46
2.5.1 Examining Effects by Headquarters Location .....	47
2.5.2 Examining Effects by Target Applicant Pro-Diversity Preferences .....	48
2.5.3 Examining Effects by Geographic Scope of Recruiting .....	49
2.5.4 Examining Intrafirm Geographic Variation .....	50
2.5.5 The Delayed Effect .....	50
2.5.6 Significant Non-Adoption.....	51
2.5.7 Examining Equal Opportunity Claims .....	52
2.5.8 Mentioning Charlottesville in Job Postings .....	52
2.6 Supplementary Wage Analysis .....	53
2.7 Discussion and Conclusion .....	54
2.7.1 Strategic Sociopolitical Positioning .....	54
2.7.2 Human Capital Strategy .....	56
2.7.3 Diversity Strategy .....	57
2.7.4 Locational Stigma and Implications for Integrated Strategy .....	58
2.7.5 A Note on Generalizability .....	59
2.8 Works Cited .....	61
2.9 Tables .....	71
2.10 Figures.....	75
2.11 Appendices.....	77
2.11.1 Methodology for Recruiting Interviewees .....	95
2.11.2 Methodology for Identifying Pro-Diversity Claims .....	96
Chapter 3 Workplace Backlash? Workforce Diversity, Status Threat, and the Contractionary Effects of Pro-Diversity Claims.....	101
3.1 Introduction.....	102
3.2 Theory .....	107
3.2.1 Prior Research on the Applicant Pool Effects of Pro-Diversity Claims .....	108
3.2.2 Anti-Diversity Backlash .....	109
3.2.3 Workplace Backlash?.....	111
3.2.4 Social Class as a Key Boundary Condition .....	112
3.3 Empirical Context .....	114
3.4 Survey Experiment .....	115
3.4.1 Target Sample .....	116
3.4.2 Experimental Procedure.....	116
3.4.3 Measures .....	119
3.5 Results.....	119
3.5.1 Summary Statistics and Balance Tests .....	119
3.5.2 Pre-Registered Analysis.....	120
3.6 Exploratory Analysis .....	122

3.6.1 Examining Alternative Mechanisms .....	122
3.6.2 Examining External Validity with Archival Data.....	125
3.7 Implications for Applicant Pool Size and Composition .....	127
3.7.1 Basic Framework .....	127
3.7.2 Assuming No Backlash.....	128
3.7.3 Allowing for Backlash .....	129
3.8 Additional Implications .....	130
3.9 Conclusion .....	132
3.10 Works Cited .....	133
3.11 Tables .....	140
3.12 Figures.....	142
3.13 Appendices.....	145
3.13.1 Qualtrics Methodology for Removing Low-Quality Respondents .....	159
Chapter 4 The Effect of Flatter Hierarchy on Applicant Pool Gender Diversity: Evidence from Experiments .....	161
4.1 Introduction.....	162
4.2 Background literature .....	166
4.2.1 Labor market sorting by gender.....	166
4.2.2 Formal hierarchy .....	167
4.2.3 How perceptions of hierarchy may affect the size of the applicant pool .....	169
4.2.4 How perceptions of hierarchy may affect the gender composition of the applicant pool .....	170
4.2.5 Theoretical summary .....	173
4.3 Methodological approach.....	173
4.4 Study 1: Field experiment.....	173
4.4.1 Sampling .....	173
4.4.2 Experimental manipulation .....	174
4.4.3 Measurement.....	176
4.4.4 Summary statistics and randomization balance.....	177
4.4.5 Estimation .....	178
4.4.6 Results .....	178
4.4.7 Post-hoc analyses of heterogeneity across job positions .....	180
4.5 Study 2: Survey experiment.....	182
4.5.1 Sampling .....	183
4.5.2 Experimental manipulations.....	184
4.5.3 Procedure.....	185
4.5.4 Measurement.....	185
4.5.5 Summary statistics and randomization balance.....	186
4.5.6 Estimation .....	187
4.5.7 Main results .....	189
4.5.8 Post-hoc mechanism analyses.....	191

4.6 Discussion .....	192
4.6.1 Theoretical contributions .....	193
4.6.2 Limitations and future research .....	195
4.7 Conclusion .....	197
4.8 Works Cited .....	198
4.9 Tables .....	205
4.10 Figures.....	210
4.11 Appendices.....	213
4.11.1 Heterogenous effects of flatter hierarchy and gender in Study 1.....	227
4.11.2 Heterogenous effects of flatter hierarchy and gender in Study 1.....	228



## List of Tables

Table 2.9.1 Summary Statistics .....	71
Table 2.9.2 Difference-in-Differences Estimates at the Level of Job Postings .....	72
Table 2.9.3 Alternative Explanations and Accompanying Empirical Analysis.....	72
Table 2.9.4 Adoption of Pro-Diversity Claims Following the Rally .....	73
Table 2.9.5 Additional Analyses Testing Theorized Mechanism.....	73
Table 2.9.6 Relationship Between Unite the Right Rally and Newfound Wage Premiums in Job Postings.....	74
Table 2.11.1 Estimated effects using employer, rather than employer-county, fixed effects.....	77
Table 2.11.2 Summary Statistics for Balanced-by-Year Panel of Employee-County Dyads.....	77
Table 2.11.3 Difference-in-differences estimates at the employer-county level show a significant increase in pro-diversity claims among Charlottesville job postings following the Unite the Right rally.....	78
Table 2.11.4 Analyzing at the employer level, the estimated effect grew in the post-treatment period. ....	78
Table 2.11.5 Consistent with the proposed mechanism, employer-level analysis shows that adoption of pro-diversity claims was larger where targeted job seekers were more likely to generalize the rally and held stronger pro-diversity preferences. ....	79
Table 2.11.6 Additional employer-level analysis is consistent with the proposed mechanism....	79
Table 2.11.7 The synthetic control method achieves better balance on observables compared to the full sample of MSAs. ....	80
Table 2.11.8 Excluding postings by UVA, the rally was associated with a newfound wage premium in Charlottesville job postings, but this premium was lower among job postings that featured pro-diversity claims. ....	80
Table 3.11.1 Subjects reported attraction to a firm that made a pro-diversity claim, a non-discrimination claim, or neither of these claims. ....	140

Table 3.11.2 Among non-bachelor's-holding subjects, the white minority information significantly decreased attraction to the pro-diversity claim, generating a newfound preference for the non-discrimination claim.....	141
Table 3.13.1 Summary of subject characteristics by treatment condition shows that randomization was successful.....	145
Table 3.13.2 Comparison of mean organizational attraction by treatment condition.....	146
Table 3.13.3 The results reported in Table A14 are unaffected by the inclusion of demographic controls.....	147
Table 3.13.4 Subjects were more likely to agree or strongly agree that the pro-diversity firm, compared to the non-discrimination firm, was more likely that other companies to employ racial minorities. This perception did not change with exposure to information. ....	148
Table 3.13.5 Subjects were not more likely to agree or strongly agree that the pro-diversity firm, compared to the non-discrimination firm, was more likely more likely than other companies to have a boss who was a racial minority. Information did not change this. ....	149
Table 3.13.6 Subjects were more likely to believe that the majority of employees within the pro-diversity firm, compared to the non-discrimination firm, were Democrats (rather than Republicans or Independents). This perception did not change with information. ....	150
Table 3.13.7 Subjects were more likely to believe that the boss of the pro-diversity firm, compared to the boss of the non-discrimination firm, was a Democrat (rather than a Republican or Independent). This perception did not change with information.....	151
Table 3.13.8 Subjects did not perceive the pro-diversity firm as more or less likely than the non-discrimination firm to unfairly discriminate against white workers. Information did not change this. ....	152
Table 3.13.9 Subjects did not perceive the pro-diversity firm as more or less likely than the non-discrimination firm to unfairly favor racial minority workers. Information did not change this.....	153
Table 3.13.10 Subjects did not perceive the pro-diversity firm as a fairer place to work than the non-discrimination firm. Information did not change this.....	154
Table 3.13.11 Non-bachelor's-holding subjects residing in labor markets that experienced greater percentage increases in racial diversity from 2010 to 2-10 were significantly less attracted to the pro-diversity firm. ....	155
Table 3.13.12 Additional test show the robustness of the finding that non-bachelor's-holding subjects residing in labor markets that experienced greater percentage increases in the share of residents of color between 2010 & 2019 were less attracted to pro-diversity.....	156
Table 3.13.13 Analysis of performance on an optional editing task. ....	157

Table 3.13.14 Analysis of treatment effects by political partisanship.....	158
Table 4.9.1 Summary statistics for Study 1.....	205
Table 4.9.2 Results using OLS regression in Study 1. Panel (a) employs the full sample, while panels (b) and (c) represent splits sample analysis by the job position that was being advertised.....	207
Table 4.9.3 Summary statistics in Study 2.....	208
Table 4.9.4 Main Results in Study 2.....	209
Table 4.10.1 Featuring a flatter hierarchical structure in recruiting efforts.....	210
Table 4.10.2 Bar graphs with 95% confidence intervals for each dependent variable in Study 1.....	211
Table 4.10.3 Main bar graphs with 95% confidence intervals for Study 2.....	212
Table 4.11.1 Correlation matrix in Study 1.....	213
Table 4.11.2 Results using logistic regression in Study 1.....	214
Table 4.11.3 Results for various thresholds for Gender API's gender prediction in Study 1.....	215
Table 4.11.4 Heterogenous effects of flatter hierarchy and gender in Study 1.....	216
Table 4.11.5 Demographic attribues of MTurk workers in Study 2.....	217
Table 4.11.6 Measurement of primary variables in Study 2.....	218
Table 4.11.7 Measurement of mechanisms in Study 2.....	219
Table 4.11.8 Correlation matrix in Study 2.....	220
Table 4.11.9 Post-hoc mechanism tests for Study 2.....	221
Table 4.11.10 Results for other potential mechanisms in Study 2.....	222

## List of Figures

Figure 2.10.1 Association Between Charlottesville and White Supremacy After the Unite the Right Rally .....	75
Figure 2.10.2 Share of Charlottesville Job Postings Featuring Pro-Diversity Claims.....	75
Figure 2.10.3 Share of Job Postings in Which Charlottesville Employers Featured Pro-Diversity Claims .....	76
Figure 2.10.4 Dynamic Difference-in-Differences Estimates of $\delta_k$ from Equation 2 .....	76
Figure 2.11.1 A photo from the rally illustrates that pro-diversity claims represented a sociopolitical stance in opposition to the anti-diversity positions of white supremacists. ....	81
Figure 2.11.2 Among employer-county dyads that feature a pro-diversity claims in any of their job postings in a given year, the percentage of employer-county dyads that feature pro-diversity claims in all of their job postings never rose above 30%. ....	82
Figure 2.11.3 Following the rally, the Charlottesville MSA rose to the 99th percentile of all MSAs in terms of job postings featuring pro-diversity claims. ....	83
Figure 2.11.4 The share of job postings made by the University of Virginia that featured pro-diversity claims increased significantly following the rally. ....	84
Figure 2.11.5 Treatment effect at the true treatment threshold is significantly larger than treatment effects estimates at placebo thresholds six months, one year, and eighteen months before treatment. ....	85
Figure 2.11.6 Excluding postings by UVA, dynamic difference-in-difference estimates show the estimated effect grew throughout the post-treatment period. ....	86
Figure 2.11.7 Robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 50% power level. ....	87
Figure 2.11.8 Robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 80% power level. ....	88
Figure 2.11.9 Excluding postings by UVA, robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 50% power level. ....	89

Figure 2.11.10 Excluding postings by UVA, robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 80% power level. ....	90
Figure 2.11.11 The difference in the share of job postings that included pro-diversity claims in the Charlottesville MSA versus the synthetic Charlottesville MSA grew significantly following the rally. ....	91
Figure 2.11.12 The difference in the share of job postings that included pro-diversity claims in the Charlottesville MSA versus the synthetic Charlottesville MSA grew significantly following the rally. ....	92
Figure 2.11.13 These lines illustrate the difference in the share of job postings featuring pro-diversity claims in the actual versus synthetic geographies for all MSAs where RMSPE equal or less than Charlottesville. Charlottesville shown in red. ....	93
Figure 2.11.14 The distribution of the ratio of pre- versus post-treatment root mean squared prediction error for all MSAs demonstrates it is very unlikely the observed effect is due to chance. ....	94
Figure 3.12.1 Subjects were randomly assigned to one of these six treatment conditions. ....	142
Figure 3.12.2 Among non-bachelor's holding subjects, the white minority information significantly decreased attraction to the pro-diversity firm. ....	142
Figure 3.12.3 For non-bachelor's-holding subjects, larger increases in racial diversity in subjects' local labor markets were associated with a preference bifurcation in favor of non-discrimination claims. ....	143
Figure 3.12.4 Without backlash, growing labor market diversity causes the pro-diversity claim to have an ever-larger expansionary effect on the applicant pool. The pro-diversity claim always results in greater applicant pool diversity. ....	143
Figure 3.12.5 With backlash, growing labor market diversity often causes the pro-diversity claim to have a contractionary effect on the applicant pool. Backlash unambiguously enhances the diversifying advantages of the pro-diversity claim. ....	144
Figure 4.11.1 Email invitation to job candidates in Study 1. ....	223
Figure 4.11.2 Recruiting material in Study 2 with no mention of the organization's hierarchical structure or the founder/CEO's gender. ....	223
Figure 4.11.3 The interventions in Study 2 inserted in the "About Us" section for the organization's hierarchical structure or the founder/CEO's gender. ....	224
Figure 4.11.4 Bar Graphs with 95% confidence intervals for mechanisms in Study 2. ....	225

Figure 4.11.5 Bar graphs with 95% confidence intervals for mechanisms by gender in Study  
2..... 226

## **Abstract**

In this dissertation, I elucidate (1) tension associated with employers' efforts to diversify applicant pools and (2) strategies whereby firms navigate this tension. In the first chapter, I present a novel theory as to why firms position on divisive, diversity-related social issues. In the second chapter, I examine whether growing labor market diversity reduces the extent to which workers from historically dominant groups are drawn to employers that make pro-diversity claims. In the final chapter, I examine whether flatter organizational structures enhance or reduce applicant pool diversity. Throughout, I use field, survey, and natural/quasi-experimental methods, which provide grounds for credible causal inference. I also outline an emergent stream of research in which I am examining the extent, origins, and implications of political de-diversification (*i.e.*, sorting by political partisanship) in the United States labor market.

## Chapter 1 Brief Introduction and Overview

Workers are a primary source of competitive advantage for firms (Becker 1964, Coff 1997). Central to my dissertation is the idea that attracting workers requires not only offering sufficiently high wages, but also successfully appealing to prospective applicants' social identities (Akerlof and Kranton 2000, Oh 2023), including race (Leibbrandt and List 2018, Flory et al. 2019), gender (Abraham and Burbano 2021, Castilla and Rho 2023), and/or political partisanship (McConnell et al. 2018). In my dissertation I elucidate strategic tension that arises as firms, in their efforts to appeal to workers from certain social groups, risk alienating those from other groups. This includes (1) work aimed at identifying this tension by demonstrating workers' diverging reactions to certain diversity-related recruitment strategies, as well as (2) work aimed at measuring how firms navigate this tension.

Although each of my dissertation chapters makes a distinct *theoretical* contribution around the unifying theme of diversity, polarization, and recruiting, my work is strongly shaped by two, closely related *phenomena* in the contemporary United States: (1) Growing ethnic diversity and (2) increasing political polarization. These trends are well-documented elsewhere (see Gest 2022 and Klein 2020). Here, I summarize them only briefly. In terms of demographic diversity, the United States is “on the precipice” of becoming a “majority-minority” nation (Craig and Richeson 2014). Whereas non-Hispanic whites have comprised a significant majority of the U.S population in recent centuries, that share is dropping precipitously and will fall below 50% in the coming



decade. In terms of polarization, recent decades have witnessed a rise in “affective polarization,” meaning dislike and distrust of those from opposing political parties (Iyengar et al. 2019). Plausibly because of this, sorting by political partisanship appears to have increased both geographically (Browns and Enos 2021, Brown et al. 2023) and along a wide range of other social dimensions (Mason 2015, Shafranek 2021, Huber and Malhorta 2017). Of course, demographic characteristics do not perfectly correspond to unified partisan identities or categories. Nevertheless, growing diversity and polarization appear related, with fear of growing diversity generating reactionary, anti-diversity political attitudes, which, in turn, motivate politicians and parties who appeal to these attitudes (Norris and Inglehart 2019, Kaufman 2019, Boxell 2020).

In my dissertation, I address questions around diversity, polarization, and recruiting that arise from these phenomena, but are not easily addressed with existing theories. For example, when and why do firms take positions on divisive, diversity-related social issues? Might growing demographic diversity shape how potential applicants respond to diversity appeals? Under pressure to cultivate greater diversity in their workforces, should employers highlight “flat” organizational structures in their recruiting efforts? Or, as I consider in ongoing work that has emerged from my dissertation, what are the degree, origins, and consequences of political de-diversification (*i.e.*, sorting by political partisanship) across workplaces?

In addressing these questions, I draw on research in political science and sociology. More so than work in strategy and management, scholarship in these fields has examined not only the implications of growing diversity (Abrajano and Hajnal 2017, Jardina 2019) and political polarization (Iyengar et al. 2019), but also deeper theoretical questions around identity-based strategies, generally *electoral* strategies, for building coalitions in divided societies (Horowitz 1985, Posner 2005, Gubler and Selway 2012, Hurst 2023). But while these disciplines have

examined wide-ranging outcomes, including intergroup attitudes (Hopkins 2010, Abascal 2020), policy preferences (Wetts and Willer 2019, Craig and Richeson 2014), voting behavior (Enos 2016), and electoral strategy (Abrajano and Hajnal 2017, Norris and Inglehart 2019, Hurst 2023), there remains a significant lacuna regarding possible effects of these phenomena on labor market preferences, and, particularly the *recruitment* strategies whereby firms appeal to these preferences. My dissertation, and, I hope, my career over the coming decades, will address this gap.

Here, I briefly summarize each of my three dissertation chapters, demonstrating how they relate to and complement one another. I conclude by outlining future work, some of it quite far along, with which I plan to move this agenda forward.

### **1.1 Chapter 1: Countervailing Claims Pro-Diversity Responses to Stigma by Association Following the Unite the Right Rally**

In my first dissertation chapter, I ask why do firms take positions on divisive social issues? The motivation here is an empirical puzzle. Given arguments that, from a strategic perspective, firms should abstain from weighing in political issues (Friedman 1970, Hadani and Schuler 2013, Burbano 2021, and Hou and Poliquin 2022), why do we observe them doing so (Larcker, Miles, and tayan 2018)? Except for a pair of theoretical papers published around the time of this chapter’s publication, prior work has deemed this positioning as non-strategic, resulting primarily from (1) top-managers simply expressing their personal beliefs (Wowak, Busenbark, and Hambrick 2022, Hambrick and Wowak 2021) or (2) isomorphic (*i.e.*, imitative) processes among peer firms (Briscoe et al. 2015).

My argument builds from the observation that the benefits firms derive from staying silent on social issues result from “false consensus” or the idea that individuals tend to presume that the unknow beliefs or preferences of others match their own (Rosenhan and Messick, 1966, Ross, Greene, and House, 1977). This argument has been leveraged in recent literature to illustrate why

political stances seem to elicit asymmetrically negative responses from those who disagree with the position (Burbano 2021, and Hou and Poliquin 2022). This research is consistent with closely related research in political science (Conover and Feldman, 1982, Tomz and Van Houweling, 2009). In fact, in my master’s thesis, which I published while writing my dissertation, I argue that a candidate’s silence regarding her social identity allows her to build larger coalitions by avoiding the censure she would otherwise face from out-group voters (Hurst 2023).

I draw on theories of stigma by association (Jonsson, Greve, and Fujiwara-Greve, 2009, Pontikes, Negro, and Rao, 2010, McDonnell, Odziemkowska, and Pontikes, 2021) to delineate conditions under which these optimistic presumptions of sociopolitical alignment are replaced by pessimistic presumptions of sociopolitical misalignment. Central to my theory is the notion that stakeholders frequently generalize the positions of non-firm political actors (such as voters, politicians, legislatures, courts, or protesters) to firms they associate with these actors, often due merely to the physical proximity of these political actors and firms. Accordingly, firms associated with political actors viewed unfavorably by the firms’ stakeholders face incentives to eschew silence and combat this stigma by claiming sociopolitical positions aligned with the preferences of the stakeholders on whom they most rely. I operationalize this in the case of Charlottesville-area employers’ pro-diversity responses to the Unite the Right Rally, which exogenously imbued these employers with an anti-diversity stigma by association.

Diversity strategy, in this first chapter, does not feature as a theoretical construct in and of itself, but rather as a way of operationalizing the broader construct of a “divisive social issue.” Similarly, recruitment serves only as the empirical context in which I operationalize the theoretical construct of “stance taking.” But although this paper is framed at the highest, most abstract level, it captures the core theme of my dissertation and the flavor of research questions I aim to pursue

in my future research agenda: Diversity as a divisive issue which firms carefully navigate in the context of recruiting.

## **1.2 Chapter 2: Workplace Backlash? Workforce Diversity, Status Threat, and the Contractionary Effects of Pro-Diversity Claims**

In my second dissertation chapter I ask: Does growing labor market diversity reduce the extent to which workers from historically dominant groups are attracted to employers that make pro-diversity claims? Unlike my first chapter, where my focus was squarely on firm behavior, I chiefly examine worker preferences and then consider their implications for firm strategy. Here, my motivation is simple: Given the extensive and ever-growing literature in political science and sociology examining how growing diversity (generally the growth in the mixed race and Hispanic populations in the United States or Muslim/African immigrant populations in Europe) shapes the *political* preferences and behaviors of historically dominant groups (generally whites in the United States and Europe), might there be analogous effects on these individuals' *labor market* preferences and behaviors? Stated differently, given evidence that growing diversity appears to be reshaping the political sphere in many Western democracies, including contributing to rise of affective polarization and reactionary parties (Kaufman 2019), might these forces come to shape labor market preferences in analogous way? Perhaps more so than my first chapter, this paper is very much a transposition of political science literature and theory to a theoretical question that lies closer to the interests of strategy and management scholars.

My focus on pro-diversity claims in recruiting is motivated by their ubiquity (in my first chapter I document that ~12% of job postings in the United States feature pro-diversity claims), and their similarity to the sort of pro-diversity public policies (such as affirmative action, open borders, and greater social spending) that appear to lose support as a result of “backlash” reactions from historically dominant groups (Wetts and Willer 2018). I draw contrast with non-

discrimination claims, which, likely due to their “colorblind” connotation, are perceived to support rather than challenge the incumbent social order (Knowles et al. 2009, Wilkins and Kaiser 2014, Leibbrandt and List 2018). Citing Du Bois’s argument that working-class whites are particularly sensitive to threats to racial status (Du Bois 1935), and subsequent work verifying and elucidating this perspective (Manstead 2018, Jardina 2019), I add that the effect should be largest among working-class members of historically dominant groups.

I examine this question with a pre-registered survey experiment of ~2,800 whites in the United States. I manipulate perceptions of labor market diversity by manipulating the length of the x-axis on chart of true demographic trends, a strategy borrowed from Abascal 2020, and then examine how exposure to this information shapes’ subjects’ reactions to language from a job posting that makes a pro-diversity claim, a non-discrimination claim, or neither of these claims. In both the main experiment and in a follow-up exploratory exercise where I examine effects relative to actual demographic change in subjects’ area of residence, I find that the effect exists *only* among whites that do not hold a bachelor’s degree.

I frame the paper around the strategic tension that my findings elucidate. Whereas prior research suggests that, as labor markets grow more diverse, pro-diversity claims will unambiguously increase applicant pool size, my results indicate conditions under which, as labor markets grow more diverse, pro-diversity claims may have a contractionary effect on applicant pool size. Across the labor market, these strategies could accelerate labor market sorting as the growing number of workers from historically marginalized groups sort into pro-diversity employers while the shrinking, but still large, share of workers from historically dominant groups eschew these employers. I illustrate this tension with a simple formal framework.

### **1.3 Chapter 3: The Effect of Flatter Hierarchy on Applicant Pool Gender Diversity: Evidence from Experiments**

My final dissertation chapter, co-authored with Justin Frake and Ronnie Lee, likewise considers unintended consequences of firms' efforts to diversify their applicant pools. In contrast to my first two dissertation chapters, we do not consider explicit pro-diversity claims. Instead, we consider how applicants' perceptions of an employer's hierarchical structure, a fundamental characteristic of all formal organizations (Burton and Obel 2004, Puranam 2018), shapes the size and gender diversity of the applicant pool.

Phenomenologically, we were motivated by the observation that, in the wake of the #MeToo movement, George Floyd's murder, and associated pressure for workforce diversity, some entrepreneurs seemed to tout their "flat" organizations structures to signal a more egalitarian alternative to the sorts of "tall," traditional hierarchies characteristic of established firms and an unjust, discriminatory status quo. In the paper, we document systematically that employers do in fact increasingly characterize their organizational structures as "flat" in their recruiting efforts. This seems especially pronounced among the types of firms, such as high-growth tech startups, in which women are chronically underrepresented (Fernandez and Campero 2017, Murciano-Goroff 2021, Phillips 2005). Rather than simply attributing this pattern to discrimination against women, which is inconsistent with evidence that women are often in highest demand where they are least represented (Lambrecht and Tucker 2019), we start from the perspective that sorting by gender might arise from differences in how women and men perceive different organizational characteristics (Barbulescu and Bidwell 2013, Fernandez-Mateo and Kaplan 2018). In other words, we take the perspective that understanding whether flatter firms achieve more gender-diverse applicants pools requires understanding differences in how they are perceived by women and men applicants.

Theoretically, we were motivated by the observation that, although extensive work has

examined the effect of hierarchical structure on current employees, essentially no work had examined the effect of hierarchy on attracting workers. Synthesizing literatures in organizational design, feminist theory, and gendered organizations, we develop divergent perspectives regarding how women and men might differ in their perceptions of organizational structure. On one hand, flatter structures might represent an “egalitarian alternative” to hierarchies which have served as an “important location of male dominance” (Acker 1990) with “masculine principles dominating their authority structures” (Kanter 1977). In this case, women may be relatively more attracted to flatter employers. On the other hand, flatter structures may connote the sort of informal work environments in which male-dominant power dynamics are able to operate unabated (Kanter 1977, Chang 2018). Here, women may perceive flatter hierarchies as signaling a “bro-culture” in which they are, relative to men, likely to face greater barriers to career advancement. In this case, men may be relatively more attracted to flatter employers.

In a large-scale, pre-registered field experiment with a partner company we find results consistent with this latter perspective: flatter structures reduce the share of women in the applicant pool. In a follow-up survey experiment we replicate this main finding, and document evidence consistent with mechanisms we theorize in this latter perspective. Compared to men, women perceive the flatter firm as less amenable to career progression, more likely to burden them with heavy workloads, and more difficult for them to fit into. In other words, it seems that women applicants may be demonstrating a preference for egalitarianism by, relative to men, avoiding flatter employers.

Our paper points to a previously undocumented tension wherein employers, in adopting certain organizational structures and/or deciding to feature these structures in their applicant pools, may risk de-diversifying their applicant pool, and by implication, their employee populations.

Although political polarization does not feature explicitly in this final chapter, we consider our question in a broader social context in which issues of gender diversity, like those of racial equality, have become highly politicized.

#### **1.4 Looking Forward: Partisan De-diversification in the workplace**

Moving forward, I am pursuing projects grounded in these dissertation chapters. Broadly stated, this work fits into an emergent agenda in which I am examining the extent, origins, and implications of partisan de-diversification (*i.e.*, sorting by political partisanship) across workplaces in the United States. Examining *extent*, my co-authors and I have a working paper in which we provide the first large-scale estimate of the degree of sorting by political partisanship in the United States labor market. Merging voter registration data with over 16 million online employee profiles covering 14 million unique workers, we find evidence of sorting, showing that compared to Republicans, Democrats work in places in which the share of co-workers that are Democrats is 12 percentage points greater, and vice versa. After accounting for sorting by geography, industry, and occupation, our estimate of partisan workplace sorting is similar in magnitude to workplace sorting by race and gender. We show, moreover, that Democrats are asymmetrically underexposed to Republican workers, but that, overall, the workplace continues to be a rare, and potentially critical, locus for the sorts of cross-partisan contact that may ameliorate affective polarization (Mutz and Mondak 2006, Wojcieszak and Warner 2020).

I am examining the *origins* of these patterns in a pair of early-stage projects. In the first, my co-authors and I examine if and how firms' diversity claims while recruiting shape the political partisan composition of the applicant pool. We have partnered with a hiring company to run a field experiment to test this relationship. In the second, my co-author and I are undertaking a replication and (significant) extension of Berniss and McDonald's (2018) finding that workers that are



political minorities (for example Democrats in majority Republican firms) are more likely to leave the firm than those in the political majority. Whereas they showed this pattern among a small sample of employees in the private equity industry, we will examine the relationship using the dataset from the working paper described above, considering approximately 14 million workers across ~450,000 employers.

In a final set of projects, my co-authors and I aim to examine the *implications* of workplace sorting by political partisanship. First, we revisit literature in management that has sought to show the relationship between workplace partisanship and firm-level outcomes by operationalizing employee partisanship with the Federal Electoral Commission's (FEC's) political donations data. We will benchmark this literature's measure against our novel measure of employee partisanship that instead uses voter registration files. We will then seek to replicate this literature's main findings using our (we believe) more informative metric of employee partisanship. In future work, I hope to examine how working in a politically diverse workplace shapes workers' political attitudes and opinions. Using this new metric, in a separate paper we will seek to examine if and how partisan diversity within firms relates to firm productivity, innovation, and performance. Here, we are motivated by extent research around gender and racial diversity including Becker's (1957) logic of discrimination (wherein employers with a taste to discriminate lose out on high-quality workers) as well as more recent arguments that diversity engenders greater range of ideas and perspectives (Ostergaard et al. 2011).

## 1.5 Works Cited

- Abascal, Maria. "Contraction as a Response to Group Threat: Demographic Decline and Whites' Classification of People Who Are Ambiguously White." *American Sociological Review* 85, no. 2 (April 1, 2020): 298–322. <https://doi.org/10.1177/0003122420905127>.
- Abraham, Mabel, and Vanessa Burbano. "Congruence Between Leadership Gender and Organizational Claims Affects the Gender Composition of the Applicant Pool: Field Experimental Evidence." *Organization Science*, March 17, 2021, orsc.2021.1442. <https://doi.org/10.1287/orsc.2021.1442>.
- Abrajano, Marisa, and Zoltan L. Hajnal. *White Backlash: Immigration, Race, and American Politics*. Princeton University Press, 2017.
- Acker, Joan. "Hierarchies, Jobs, Bodies:: A Theory of Gendered Organizations." *Gender & Society* 4, no. 2 (June 1, 1990): 139–58. <https://doi.org/10.1177/089124390004002002>.
- Akerlof, George A., and Rachel E. Kranton. "Economics and Identity." *The Quarterly Journal of Economics* 115, no. 3 (August 1, 2000): 715–53. <https://doi.org/10.1162/003355300554881>.
- Barbulescu, Roxana, and Matthew Bidwell. "Do Women Choose Different Jobs from Men? Mechanisms of Application Segregation in the Market for Managerial Workers." *Organization Science* 24, no. 3 (June 2013): 737–56. <https://doi.org/10.1287/orsc.1120.0757>.
- Becker, Gary. *The Economics of Discrimination*. University Of Chicago Press, 1957.
- Becker, Gary S. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. University of Chicago Press, 2009.
- Bermiss, Y. Sekou, and Rory McDonald. "Ideological Misfit? Political Affiliation and Employee Departure in the Private-Equity Industry." *Academy of Management Journal*, February 27, 2018. <https://doi.org/10.5465/amj.2016.0817>.
- Boxell, Levi. "Demographic Change and Political Polarization in the United States." *Economics Letters* 192 (July 1, 2020): 109187. <https://doi.org/10.1016/j.econlet.2020.109187>.
- Brown, Jacob R, and Enrico Cantoni. "The Increase in Partisan Segregation in the United States," 2023.
- Brown, Jacob R., and Ryan D. Enos. "The Measurement of Partisan Sorting for 180 Million Voters." *Nature Human Behaviour* 5, no. 8 (August 2021): 998–1008. <https://doi.org/10.1038/s41562-021-01066-z>.

- Burbano, Vanessa C. “Social Responsibility Messages and Worker Wage Requirements: Field Experimental Evidence from Online Labor Marketplaces.” *Organization Science* 27, no. 4 (August 2016): 1010–28. <https://doi.org/10.1287/orsc.2016.1066>.
- Burton, Richard M., and Borge Obel. *Strategic Organizational Diagnosis and Design: The Dynamics of Fit*. Springer Science & Business Media, 2004.
- Castilla, Emilio J., and Hye Jin Rho. “The Gendering of Job Postings in the Online Recruitment Process.” *Management Science* 69, no. 11 (November 2023): 6912–39. <https://doi.org/10.1287/mnsc.2023.4674>.
- Chang, Emily. *Brotopia: Breaking Up the Boys’ Club of Silicon Valley*. Penguin, 2018.
- Coff, Russell W. “Human Assets and Management Dilemmas: Coping with Hazards on the Road to Resource-Based Theory.” *The Academy of Management Review* 22, no. 2 (1997): 374–402. <https://doi.org/10.2307/259327>.
- Conover, Pamela Johnston, and Stanley Feldman. “Projection and the Perception of Candidates’ Issue Positions.” *The Western Political Quarterly* 35, no. 2 (1982): 228–44. <https://doi.org/10.2307/448017>.
- Craig, Maureen A., and Jennifer A. Richeson. “On the Precipice of a ‘Majority-Minority’ America: Perceived Status Threat From the Racial Demographic Shift Affects White Americans’ Political Ideology.” *Psychological Science* 25, no. 6 (June 1, 2014): 1189–97. <https://doi.org/10.1177/0956797614527113>.
- Du Bois, W. E. B. *Black Reconstruction in America 1860-1880*. Simon and Schuster, 1935.
- Fernandez, Roberto M., and Santiago Campero. “Gender Sorting and the Glass Ceiling in High-Tech Firms.” *ILR Review* 70, no. 1 (January 1, 2017): 73–104. <https://doi.org/10.1177/0019793916668875>.
- Fernandez-Mateo, Isabel, and Sarah Kaplan. “Gender and Organization Science: Introduction to a Virtual Special Issue.” *Organization Science* 29, no. 6 (December 2018): 1229–36. <https://doi.org/10.1287/orsc.2018.1249>.
- Flory, Jeffrey A., Andreas Leibbrandt, Christina Rott, and Olga Stoddard. “Increasing Workplace Diversity: Evidence from a Recruiting Experiment at a Fortune 500 Company.” *Journal of Human Resources*, August 6, 2019, 0518. <https://doi.org/10.3368/jhr.56.1.0518-9489R1>.
- Friedman, Milton. “A Friedman Doctrine-- The Social Responsibility Of Business Is to Increase Its Profits.” *The New York Times*, September 13, 1970, sec. Archives. <https://www.nytimes.com/1970/09/13/archives/a-friedman-doctrine-the-social-responsibility-of-business-is-to.html>.

- Gubler, Joshua R., and Joel Sawat Selway. "Horizontal Inequality, Crosscutting Cleavages, and Civil War." *Journal of Conflict Resolution* 56, no. 2 (April 1, 2012): 206–32. <https://doi.org/10.1177/0022002711431416>.
- Hadani, Michael, and Douglas Schuler. "In Search of El Dorado: The Elusive Financial Returns on Corporate Political Investments." *Strategic Management Journal* 34, no. 2 (2013): 165–81.
- Hambrick, Donald C., and Adam J. Wowak. "CEO Sociopolitical Activism: A Stakeholder Alignment Model." *Academy of Management Review* 46, no. 1 (January 2021): 33–59. <https://doi.org/10.5465/amr.2018.0084>.
- Hopkins, Daniel J. "Politicized Places: Explaining Where and When Immigrants Provoke Local Opposition." *The American Political Science Review* 104, no. 1 (2010): 40–60.
- Horowitz, Donald L. *Ethnic Groups in Conflict*. University of California Press, 1985.
- Hou, Young, and Christopher W. Poliquin. "The Effects of CEO Activism: Partisan Consumer Behavior and Its Duration." *Strategic Management Journal* 44, no. 3 (2023): 672–703. <https://doi.org/10.1002/smj.3451>.
- Huber, Gregory A., and Neil Malhotra. "Political Homophily in Social Relationships: Evidence from Online Dating Behavior." *The Journal of Politics* 79, no. 1 (January 2017): 269–83. <https://doi.org/10.1086/687533>.
- Hurst, Reuben. "Strategically Ambiguous Identities: Evidence from Evangelical Clergy in Brazilian Elections." *The Journal of Politics* 85, no. 4 (October 2023): 1336–51. <https://doi.org/10.1086/723969>.
- Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22, no. 1 (2019): 129–46. <https://doi.org/10.1146/annurev-polisci-051117-073034>.
- Jardina, A. *White Identity Politics*. Cambridge Studies in Public Opinion and Political Psychology. Cambridge University Press, 2019. <https://books.google.com/books?id=DUPXvAEACAAJ>.
- Jonsson, Stefan, Henrich R. Greve, and Takako Fujiwara-Greve. "Undeserved Loss: The Spread of Legitimacy Loss to Innocent Organizations in Response to Reported Corporate Deviance." *Administrative Science Quarterly* 54, no. 2 (June 1, 2009): 195–228. <https://doi.org/10.2189/asqu.2009.54.2.195>.
- Kanter, Rosabeth Moss. "Some Effects of Proportions on Group Life: Skewed Sex Ratios and Responses to Token Women." *American Journal of Sociology* 82, no. 5 (1977): 965–90.

- Kaufmann, E. *Whiteshift: Populism, Immigration and the Future of White Majorities*. Penguin Books Limited, 2018. <https://books.google.com/books?id=tFZWDwAAQBAJ>.
- Klein, Ezra. *Why We're Polarized*. Simon and Schuster, 2020.
- Knowles, Eric D., Brian S. Lowery, Caitlin M. Hogan, and Rosalind M. Chow. "On the Malleability of Ideology: Motivated Construals of Color Blindness." *Journal of Personality and Social Psychology* 96, no. 4 (April 2009): 857–69. <https://doi.org/10.1037/a0013595>.
- Larcker, David F, Stephen A Miles, and Brian Tayan. "The Double-Edged Sword of CEO Activism." *Rock Center for Corporate Governance at Stanford University Closer Look Series: Topics, Issues and Controversies in Corporate Governance No. CGRP-74*, 2018, 12.
- Leibbrandt, Andreas, and John A. List. "Do Equal Employment Opportunity Statements Backfire? Evidence From A Natural Field Experiment On Job-Entry Decisions." National Bureau of Economic Research, September 17, 2018. <https://doi.org/10.3386/w25035>.
- Manstead, Antony S. R. "The Psychology of Social Class: How Socioeconomic Status Impacts Thought, Feelings, and Behaviour." *The British Journal of Social Psychology* 57, no. 2 (April 2018): 267–91. <https://doi.org/10.1111/bjso.12251>.
- Mason, Lilliana. "'I Disrespectfully Agree': The Differential Effects of Partisan Sorting on Social and Issue Polarization." *American Journal of Political Science* 59, no. 1 (January 2015): 128–45. <https://doi.org/10.1111/ajps.12089>.
- McDonnell, Mary-Hunter, Kate Odziemkowska, and Elizabeth Pontikes. "Bad Company: Shifts in Social Activists' Tactics and Resources After Industry Crises." *Organization Science* 32, no. 4 (July 2021): 1033–55. <https://doi.org/10.1287/orsc.2020.1410>.
- Murciano-Goroff, Raviv. "Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers." *Management Science*, September 16, 2021. <https://doi.org/10.1287/mnsc.2021.4077>.
- Mutz, Diana C., and Jeffery J. Mondak. "The Workplace as a Context for Cross-Cutting Political Discourse." *The Journal of Politics* 68, no. 1 (February 2006): 140–55. <https://doi.org/10.1111/j.1468-2508.2006.00376.x>.
- Norris, Pippa, and Ronald Inglehart. *Cultural Backlash and the Rise of Populism: Trump, Brexit, and Authoritarian Populism*. Cambridge University Press, 2019.
- Oh, Suanna. "Does Identity Affect Labor Supply?" *American Economic Review* 113, no. 8 (August 2023): 2055–83. <https://doi.org/10.1257/aer.20211826>.

- Østergaard, Christian R., Bram Timmermans, and Kari Kristinsson. “Does a Different View Create Something New? The Effect of Employee Diversity on Innovation.” *Research Policy* 40, no. 3 (April 1, 2011): 500–509. <https://doi.org/10.1016/j.respol.2010.11.004>.
- Phillips, Damon J. “Organizational Genealogies and the Persistence of Gender Inequality: The Case of Silicon Valley Law Firms.” *Administrative Science Quarterly* 50, no. 3 (September 1, 2005): 440–72. <https://doi.org/10.2189/asqu.2005.50.3.440>.
- Pontikes, Elizabeth, Giacomo Negro, and Hayagreeva Rao. “Stained Red: A Study of Stigma by Association to Blacklisted Artists during the ‘Red Scare’ in Hollywood, 1945 to 1960.” *American Sociological Review* 75, no. 3 (June 1, 2010): 456–78. <https://doi.org/10.1177/0003122410368929>.
- Posner, Daniel. *Institutions and Ethnic Politics in Africa*. Political Economy of Institutions and Decisions. Cambridge University Press, 2005.
- Puranam, Phanish. *The Microstructure of Organizations*. Oxford University Press, 2018. <https://doi.org/10.1093/oso/9780199672363.001.0001>.
- Rosenhan, David, and Samuel Messick. “Affect and Expectation.” *Journal of Personality and Social Psychology* 3, no. 1 (1966): 38–44. <https://doi.org/10.1037/h0022633>.
- Ross, Lee, David Greene, and Pamela House. “The ‘False Consensus Effect’: An Egocentric Bias in Social Perception and Attribution Processes.” *Journal of Experimental Social Psychology* 13, no. 3 (May 1, 1977): 279–301. [https://doi.org/10.1016/0022-1031\(77\)90049-X](https://doi.org/10.1016/0022-1031(77)90049-X).
- Shafranek, Richard M. “Political Considerations in Nonpolitical Decisions: A Conjoint Analysis of Roommate Choice.” *Political Behavior* 43, no. 1 (March 1, 2021): 271–300. <https://doi.org/10.1007/s11109-019-09554-9>.
- Tomz, Michael, and Robert Van Houweling. “The Electoral Implications of Candidate Ambiguity.” *American Political Science Review* 103, no. 1 (February 2009): 83–98. <https://doi.org/10.1017/S0003055409090066>.
- Wetts, Rachel, and Robb Willer. “Privilege on the Precipice: Perceived Racial Status Threats Lead White Americans to Oppose Welfare Programs.” *Social Forces* 97, no. 2 (December 1, 2018): 793–822. <https://doi.org/10.1093/sf/soy046>.
- Wilkins, Clara L., and Cheryl R. Kaiser. “Racial Progress as Threat to the Status Hierarchy: Implications for Perceptions of Anti-White Bias.” *Psychological Science* 25, no. 2 (February 1, 2014): 439–46. <https://doi.org/10.1177/0956797613508412>.
- Wojcieszak, Magdalena, and Benjamin R. Warner. “Can Interparty Contact Reduce Affective Polarization? A Systematic Test of Different Forms of Intergroup Contact.” *Political*

*Communication* 37, no. 6 (November 1, 2020): 789–811.  
<https://doi.org/10.1080/10584609.2020.1760406>.

Wowak, Adam J., John R. Busenbark, and Donald C. Hambrick. “How Do Employees React When Their CEO Speaks Out? Intra- and Extra-Firm Implications of CEO Sociopolitical Activism.” *Administrative Science Quarterly*, February 17, 2022.  
<https://doi.org/10.1177/00018392221078584>.

## **Chapter 2 Countervailing Claims Pro-Diversity Responses to Stigma by Association Following the Unite the Right Rally**

### **Abstract**

Why do firms take positions on divisive social issues? In this article, I draw on theories of stigma by association to explain why firms' mere proximity to controversial political actors may lead stakeholders to presume that firms silent on social issues are misaligned with the stakeholders' sociopolitical preferences. Firms, in turn, countervail these presumptions of misalignment by eschewing silence and claiming sociopolitical positions. Substantiating this theory in the context of employee recruitment following the 2017 Unite the Right White supremacist rally in Charlottesville, Virginia, I show that Charlottesville's employers combated presumptions that they shared demonstrators' anti-diversity positions by making countervailing pro-diversity claims in their online job postings. In supplementary analysis, I show that the rally was associated with a newfound wage premium in job postings by Charlottesville's employers but that this premium was lower when employers made pro-diversity claims. This study advances understanding of strategic sociopolitical positioning whereby firms make calculated appeals to stakeholders. It contrasts with related research showing that firms use social claims to combat negative evaluations resulting from their own actions or to differentiate from competitors.



## 2.1 Introduction

Why do firms take positions on divisive social issues? Despite some evidence that stakeholders reward firms that take positions aligned with the former's preferences (Chatterji and Toffel 2019, Wowak et al. 2022), growing research suggests that firms face significant, offsetting motivations to stay silent. Not only does silence preclude accusations of hypocrisy (Wagner et al. 2009, Janney and Gove 2011, Carlos and Lewis 2018) and unwanted pressure from activists (Vogel 2005, Baron and Diermeier 2007), but it may even cultivate false consensus, wherein stakeholders with varying preferences presume that silent firms align with those respective preferences (Burbano 2021, Hou and Poliquin 2022). In this account, presumptions of sociopolitical alignment mean that sociopolitical claims do little to enhance stakeholder perceptions that firms are aligned with the former's preferences while also eliciting negative reactions from stakeholders misaligned with a stance. Closely related research in political science similarly demonstrates that silence can forestall scrutiny while allowing candidates to build larger, more-diverse coalitions (Conover and Feldman, 1982, Tomz and Van Houweling 2009, Hurst 2023).

But while this pessimistic perspective on the benefits of sociopolitical positioning largely supports “conventional wisdom” (Hambrick and Wowak 2021) that firms fare better when they remain silent (Friedman 1970, Hadani and Schuler 2013, Bhagwat et al. 2020), it raises questions about why firms often fail to do so (Larcker, Miles, and Tayan 2018). For example, in the wake of Texas's Senate Bill 8, which in 2021 codified a strongly pro-life stance on abortion, many Texas-based firms took highly publicized, pro-choice positions (Pardes 2021). In the wake of Donald Trump's victory in the 2016 presidential election, many U.S. firms made commitments to racial, gender, and LGBTQ+ diversity (*The Economist* 2017). And, as I examine in detail in this article, in the wake of a large-scale White supremacist rally in Charlottesville, Virginia, in 2017, many

Charlottesville employers added pro-diversity claims to their job postings. When firms take stances in these and other cases, do their actions primarily manifest top managers' values or result from imitative processes among peer firms (Marquis 2003, Briscoe et al. 2015)? When, in contrast, do such instances reflect calculated, strategic appeals to stakeholders (Melloni et al. 2023, Mohliver, et al. 2023)?

I draw on theories of stigma by association (Jonsson et al. 2009, Pontikes et al. 2010, McDonnell et al. 2021) to explain one reason that firms strategically take positions on sociopolitical issues. Central to my theory is the notion that stakeholders frequently generalize the positions of non-firm political actors (such as voters, politicians, legislatures, courts, or protesters) to firms they associate with these actors, often due merely to the physical proximity of these political actors and firms. Among stakeholders whose preferences do not align with these actors' positions, this generalization engenders "sociopolitical stigma," which, drawing on Devers et al. (2009), I define as stakeholders' negative evaluations of a firm due to a perceived lack of alignment between this firm and stakeholders' sociopolitical positions. By replacing stakeholders' optimistic presumptions of sociopolitical alignment with pessimistic presumptions of sociopolitical misalignment, sociopolitical stigma by association undermines the feasibility of maintaining false consensus through silence.

Accordingly, firms associated with non-firm political actors viewed unfavorably by the firms' stakeholders face incentives to eschew silence and combat this stigma by claiming sociopolitical positions aligned with their stakeholders' preferences. Drawing on Werner's (2012) theory of private politics as defensive, strategic maneuvering, I propose that firms respond to these incentives by making "countervailing claims": sociopolitical claims meant to offset or counter the negative effects of stigma by association. In my account, firms take sociopolitical positions not to

offset negative evaluations arising from their own actions (Elsbach, 1994, Bansal and Clelland, 2004, McDonnell and King, 2013) or to counterposition vis-à-vis rival firms (Mohliver, Crilly, and Kaul, 2023) but, rather, to distance themselves from controversial non-firm political actors with whom they are associated.

I substantiate this theory in the context of employee recruitment following the 2017 Unite the Right White supremacist rally in Charlottesville, Virginia. Using a dataset of 66.5 million online jobs postings, I show that in the two years following the rally, employers in Charlottesville significantly increased their use of pro-diversity claims in their job postings. Consistent with my theory, the adoption of pro-diversity claims was not exceptionally widespread but was most pronounced when prospective applicants plausibly (1) were more inclined to generalize the White supremacists' anti-diversity positions to Charlottesville's employers and (2) held stronger pro-diversity preferences. In supplementary analysis, I show that the rally was associated with a newfound wage premium but that this premium was lower when employers made pro-diversity claims. I motivate and supplement my quantitative analysis with qualitative evidence from memoirs, media reports, and 21 original field interviews.

## **2.2 Theory**

Firms increasingly take positions on divisive social issues. This positioning is distinct from the “pro-social” claim-making long examined in the impression management literature inasmuch as the former involves positions that are contested rather than “socially acceptable” (McDonnell and King, 2013: 388). Although nascent research in this area has tended to examine CEO activism (Chatterji and Toffel 2019, Hambrick and Wowak 2021, Hou and Poliquin 2022, Wowak, et al. 2022), firms' sociopolitical positioning is often not easily attributable to any single organizational member. Moreover, while some research narrowly considers position-taking with respect to issues

unrelated to a firm's business, in many cases a firm's sociopolitical positioning does relate to business, including outdoor recreation companies taking pro-environmental positions, employers taking positions on their employees' reproductive rights, or, more generally, firms with diverse workforces proclaiming the value of diversity. In contrast to corporate social responsibility or other forms of non-market strategy, sociopolitical positioning is distinct in that it can be achieved with little expenditure of company resources (Wowak et al. 2022). This is not to say that more-costly, substantive actions cannot complement sociopolitical positioning but, rather, that the act of claiming a position is itself a theoretical construct of interest. This distinction mirrors a long tradition in political science that treats positioning as theoretically distinct from actual policies undertaken to pursue these positions (Dewan and Shepsle 2011).

### 2.2.1 Stakeholder Responses to Sociopolitical Positioning

With little exception, emergent scholarly attention to firms' sociopolitical positioning has examined whether these actions benefit firms. Highlighting upsides, some research emphasizes that positioning requires little expenditure of firm resources while increasing the extent to which stakeholders aligned with the espoused position reward the firm. On the first point, firms can take positions via simple statements in social media posts, television ads, or, as I will show, job postings, compared to more-resource-intensive non-market strategies such as self-regulation (Short and Toffel, 2010), political donations (Werner 2017), corporate volunteer programs (Carnahan et al. 2017), or philanthropy (Luo et al. 2018). On the second point, workers (Carpenter and Gong 2015, McConnell et al. 2018, Wowak et al. 2022), consumers (Chatterji and Toffel, 2019, Panagopoulos et al., 2020), and even investors (Mohliwer and Hawn, 2019, Bolton et al., 2020) may be inclined to reward firms that take positions aligned with these groups' preferences.

A separate set of studies has illustrated significant, offsetting downsides of sociopolitical positioning. This research emphasizes that positioning may incur unwanted public scrutiny and asymmetrically negative reactions from stakeholders misaligned with the espoused position. Baron and Diermeier (2007) as well as Vogel (2005) suggested that stance-taking attracts activists in favor of the claimed position, who perceive the firm taking the position as sympathetic to their cause and, consequently, as less likely to resist pressure to pursue costly actions consistent with the position. For example, Starbucks's pro-environmental positioning drew attention from activists who, in turn, successfully pressured the firm to embrace costly changes to its coffee-sourcing procedures (Argenti 2004). When scrutiny increases, there is also increased risk that a firm will be seen as hypocritical, which can significantly reduce stakeholder support (Vogel 2005, Wagner et al. 2009, Janney and Gove 2011). For example, many organizations that stated support for the Black Lives Matter movement in the summer of 2020, including J. P. Morgan Chase (Jan et al. 2020), L'Oréal (Elan 2020), and the United Kingdom's Royal Academy of Drama and Art (Bakare 2020), were subsequently accused of hypocrisy for actions inconsistent with these claims. Silence on sociopolitical issues forestalls these accusations and affords latitude for actions that might otherwise be deemed hypocritical (Carlos and Lewis 2018).

Examining consumer reactions to corporate support for gun control in the wake of mass shootings, Hou and Poliquin (2022) demonstrated asymmetrically negative responses from consumers who oppose gun control. Similarly, Burbano (2021) presented evidence that taking a position on North Carolina's trans-exclusionary bathroom bill failed to increase output among workers aligned with the position but significantly reduced output among workers misaligned with the position. Burbano proposed that these findings may be driven by false consensus bias, which is the tendency to overestimate the extent to which the unknown preferences of others match one's

own (Rosenhan and Messick 1966, Ross et al. 1977). These findings are also consistent with theories of negativity bias and loss aversion, according to which bad outcomes, in this case sociopolitical misalignment, have stronger behavioral implications than good outcomes produce, in this case sociopolitical alignment (Kahneman et al. 1991, Baumeister et al. 2001, Rozin and Royzman 2001).

These findings suggest that claiming a sociopolitical position does little to improve perceptions of alignment among those aligned with the position and elicits outside negative reactions among those misaligned with the position. The idea that claiming a sociopolitical position undermines stakeholders' optimistic presumptions that the firm aligns with their preferences is consistent with related research in political science. By avoiding clear policy stances, political candidates can often attract a wider range of voters since these voters presume that the candidates' positions align with their own different preferences (Markus and Converse 1979, Conover and Feldman 1982, Tomz and Van Houweling 2009). Inasmuch as candidates take well-defined positions, they are incentivized to do so on only the few core issues on which voters expect them to take a position (Krosnick 1990). This logic is especially germane to firms. If a firm fails to communicate quarterly earnings, the goods it sells, the salaries it pays, or other information most core to business, this silence will unavoidably elicit negative attention from investors, customers, and employees. Silence on divisive issues, however, may deflect unwanted scrutiny while leaving room for stakeholders to make optimistic presumptions regarding the firm's stances on these issues. These strong motivations to remain silent perhaps explain why firms are silent on the vast majority of social issues. They also raise questions about when, if ever, decisions to eschew silence and claim positions on these issues are strategic. They suggest that firms' taking

positions on sociopolitical issues may simply reflect the values of top managers (Hambrick and Wowak 2021) or result from imitative processes among peer firms (Marquis 2003).

### 2.2.2 Combating Sociopolitical Stigma with Countervailing Claims

I propose conditions under which firms face strong incentives to eschew silence and strategically claim positions on divisive issues. My argument builds on Burbano's (2021) observation that stakeholders make presumptions about the positions of silent firms. But whereas Burbano pointed only to optimistic presumptions arising from false consensus bias, I argue that these presumptions may sometimes be pessimistic. Pessimistic presumptions of sociopolitical misalignment erode the benefits of sociopolitical silence and incentivize firms to combat these presumptions with countervailing claims. In the following two subsections, I detail this two-part process wherein (1) stakeholders generalize non-firm political actors' positions to firms associated with these actors, and consequently, (2) firms that depend on stakeholders whose preferences misalign with these positions combat this generalization with countervailing claims.

***Sociopolitical stigma.*** My argument starts from the premise that stakeholders generalize the positions of high-profile, non-firm political actors to firms associated with these actors. Frequently, this association will arise due to the firm's location. For example, stakeholders observe a state legislature pass left-leaning laws and then presume that firms in that state tend to hold left-leaning sociopolitical positions. These generalizations, however, might arise also due to associations unrelated to location. For example, stakeholders might observe White supremacist rioters bearing Tiki torches and then presume that the vendors or manufacturers of Tiki torches share the rioters' grievances. Critically, these presumptions arise from mere association and not because stakeholders see evidence that a firm's actions are consistent with these positions. When political actors misalign with stakeholders' preferences, this generalization engenders "sociopolitical

stigma,” which I define as stakeholders’ negative evaluation of a firm due to a perceived lack of alignment between this firm and stakeholders’ sociopolitical positions. This definition is a variation of Devers et al.’s (2009) definition of stigma as “a label that evokes a collective perception that [an] organization is deeply flawed and discredited.” In my construct, the discrediting “label” is distinctly sociopolitical: stakeholders with a shared sociopolitical preference collectively perceive these stigmatized firms as misaligned with their preferences and are consequently less inclined to reward them.

This argument is consistent with prior research demonstrating that individuals generalize the actions of focal entities to sets of associated entities (Adut 2005, Pontikes et al. 2010). The canonical example is Hollywood during the Red Scare, when audiences formed negative evaluations not only of actors who had been blacklisted as Communists but also of actors who had not been listed as Communists but had previously worked with blacklisted actors (Pontikes et al. 2010). This inclination to generalize is starkly illustrated by evidence that individuals can incur negative evaluations due simply to physical proximity to individuals with stigmatized characteristics (Hebl and Mannix 2003, Penny and Haddock 2007, Pryor et al. 2012). These generalization processes play a critical role in how stakeholders evaluate organizations. Unable to collect detailed information on large numbers of individual organizations, stakeholders frequently make wider inferences based on the actions of a small number of organizations. These interorganizational generalization processes have been documented in many contexts, including among organizations that share a parent organization (Jensen 2006, Piazza and Jourdan 2018), industry affiliation (Jonsson et al. 2009, Durand and Vergne 2015, Naumovska and Lavie 2021), partnership (McDonnell et al. 2021), or organizational form (Yue et al. 2013).



Through these same generalization processes, political actors' position-taking significantly shapes stakeholders' presumptions regarding the positions of firms associated with these actors. Political actors attract media attention, which coordinates generalization processes and results in "collective perceptions" of firms' sociopolitical positions, which, because large numbers of stakeholders hold them, become material to the firm (Devers et al. 2009). Other forces or actors also plausibly shape stakeholders' presumptions about firms' sociopolitical positions. The prominence of political actors in the sociopolitical landscape, however, as well as their mandates to frequently take sociopolitical positions, means that they play a primary role.

***Countervailing claims.*** These presumptions create incentives for firms to either claim a sociopolitical position or maintain silence. On the one hand, if the firm relies heavily on stakeholders whose preferences are further from these actors' positions, then the firm is more incentivized to combat the ensuing stigma through countervailing sociopolitical claims. In this case, the firm faces a greater burden of proof in demonstrating to stakeholders that, despite association with a political actor who holds a disliked position, the firm is, in fact, aligned with its stakeholders' preferences. In terms of Adut's (2005) evocative language, such firms have become, in the eyes of their stakeholders, contaminated "third parties" who, for reasons unrelated to their own actions, face pressure to demonstrate "extraordinary zeal" as they "signal rectitude or resolve" to their stakeholders' preferences. In doing so, the firm still incurs costs associated with making these claims, such as unwanted attention from activists or accusations of hypocrisy, but these costs are more likely to be offset by the benefits of remedying stakeholders' presumptions of misalignment. On the other hand, if the firm relies more heavily on stakeholders whose preferences are closer to the political actor's position, then the firm benefits from presumptions of sociopolitical alignment and thus has a greater incentive to remain silent.

My characterization of strategic positioning is consistent with Werner's (2012: 22–23) conceptualization of private politics, which modifies March's behavioral theory of the firm to allow more room for "strategic calculation." Central to both my argument and Werner's (2012: 24) conceptualization is the idea that firms monitor the social environment and seek to "protect themselves and their reputations from hostile public opinion." I argue, in contrast to this and other prior work, that sociopolitical positioning is not a direct response to competitors (Mohliver et al. 2023), an effort to assuage general public distrust of corporations, or meant to compensate for the focal firm's own missteps (Elsbach 1994, Werner 2012, McDonnell and King 2013). Instead, countervailing claims arise as decision makers within firms come to believe that due to their mere association with political actors who take positions disliked by their stakeholders, they may benefit by combating presumptions that they share these disliked positions. Note that this argument operates regardless of exactly why perceptions of sociopolitical misalignment make stakeholders less inclined to reward firms. Whether this disinclination is fundamentally taste-based (i.e., due simply to ideological aversion) or statistical (i.e., due to concern that the sociopolitical position signals additional undesirable firm qualities), firms face incentives to remedy it by explicitly aligning with the preferences of the stakeholders on whom they most depend.

### **2.3 Empirical Context: Unite the Right Rally in Charlottesville, Virginia**

I examine this theory in the context of employers' positioning on diversity issues in the wake of the Unite the Right White supremacist rally in Charlottesville, Virginia, on August 11–12, 2017. This context is ideal for two key reasons. First, the rally was exogenous to the actions of Charlottesville's employers. Although political actors constantly take positions on divisive issues, this positioning is often endogenous in firms' other actions. For example, firms donate to politicians, lobby for legislation, participate in lawsuits, and officially sponsor parades and other

protests. At the very least, they may employ individuals who privately engage in these activities. These activities make it difficult to infer whether firms' sociopolitical positioning is a response to the positions taken by political actors or is, instead, driven by some other firm characteristic. In the case of the Unite the Right rally, however, White supremacist protesters came, with almost no exception, from outside of Charlottesville, thereby plausibly staining Charlottesville-area employers with an anti-diversity stigma that was unrelated to their own actions and even the private behaviors of their employees. Second, Charlottesville is isolated in rural central Virginia. This isolation allowed me to compare the responses of employers in this clearly defined treated location to the behavior of employers in control locations. In the remainder of this section, I describe the Unite the Right rally and provide preliminary, qualitative evidence from memoirs, media reports, and 21 original field interviews that Charlottesville employers sought to combat anti-diversity stigma arising from the rally through countervailing pro-diversity claims. This illustrative, suggestive preliminary qualitative evidence serves primarily to describe the setting and build intuition for the large-data, quantitative analysis that follows.

### 2.3.1 Sociopolitical Stigma for Charlottesville Employers

The Unite the Right rally was organized in response to the Charlottesville city council's decision in February 2017 to remove a statue of Confederate Civil War general Robert E. Lee. The rally attracted thousands of White supremacists from around the country to Charlottesville on August 11 and 12, 2017. Through their clothing, protest signs, and verbal chants, the agitators manifested strong anti-diversity positions. The event captured widespread attention, with striking images of mostly male, Tiki torch-bearing White supremacists appearing across major media outlets (McAuliffe 2019, Signer 2020).

Multiple accounts describe how the rally stained Charlottesville with a newfound anti-diversity stigma. In his memoir of the events, Terry McAuliffe (2019: 158), governor of Virginia at the time, characterized the rally as a “lightning bolt” that “lit up the scourge of racism and hatred of others as it really is, in the here and now.” Then-Charlottesville mayor Michael Signer characterized the rally as an “earthquake” in which perceptions of Charlottesville were “fundamentally shaken.” Signer (2020: 8) lamented that, due to the rally,

. . . Charlottesville would become synonymous with white supremacy and with terrorism.

The city’s “brand” now comingled the pleasant contours of UVA and Monticello with swastikas, shields, swords, helmets, and that haunting chant: “Jews will not replace us.”

Some Charlottesville residents lamented a disconnect between these perceptions and their experiences. One resident quoted in *The New York Times* characterized “the portrait of Charlottesville that emerged from 2017” as “exaggerated,” bemoaning “a misconception that we just have Klan members hanging around Charlottesville” (MacFarquhar 2021). Similarly, multiple interviewees described friends and family from out of town who had difficulty recognizing that almost no rioters were from Charlottesville, and thus these observers formed an inaccurate belief that the rally displayed anti-diversity attitudes that were common among Charlottesville’s residents. One interviewee recalled that while attending college out of state, her son was introduced as coming from “Charlottesville, where the Nazis ride.” She added, “that’s just the way outsiders perceived it.”

To more systematically substantiate this newfound stigma by association, I documented media mentions of Charlottesville in the periods before and after the rally. Using the LexisNexis media archives, which capture major print, audio, and television media reports, I documented the number of stories by month from 2015 through 2019 that contained the search terms

“Charlottesville,” “White supremacy,” or both of these terms. Illustrating these patterns, Figure 1 shows strongly correlated upticks in each of these three search terms beginning in August 2017. The chart also shows correlated upticks in these search terms around the annual anniversary of the events (August 2018 and August 2019), suggesting that this newfound stigma was not transitory.

My interviewees provided preliminary evidence that Charlottesville’s employers became sensitive to how this stigma by association shaped prospective employees’ perceptions of the firms. A Charlottesville government official recalled,

I can remember specifically a conversation with a large company in town that, you know, they had visits planned and they had to change those. And then they had people kind of [*pause*] concerned. . . . I think the local firms that were hiring, you know, significantly and looking to recruit people here, to take a job, had to kind of adjust their approach and be prepared to answer the question about [where they stood on the issue of diversity].

Reporting similar concerns, another interviewee commented,

Every time I try to hire someone at the highest level we have to talk about [the rally]. If they’re from Stanford, Duke, Berkeley, they’re like what the hell was that? Is that who you are?

Other interviewees mentioned similar difficulties. One expressed that the rally “came up repeatedly in interviews” and that “one African American applicant asked me explicitly about it and how [the employer] was resisting it.”

### 2.3.2 Countervailing Pro-Diversity Claims

The Unite the Right rally put Charlottesville’s employers in a difficult situation regarding whether they should speak up or remain silent on issues of diversity. Although most Americans do not explicitly support White supremacy, there are surprisingly high levels of agreement with anti-

diversity positions central to White supremacist ideology. In a national poll conducted one month after the rally, fewer than 8 percent of respondents expressed support for “White nationalism” or “neo-Nazism,” but 31 percent agreed that the United States needed to “protect and preserve its White European heritage,” 39 percent agreed that “White people are under attack in this country,” and 57 percent said that Confederate monuments should remain in public spaces (UVA Center for Politics, 2017).

Disagreement about whether to respond to the Unite the Right rally with countervailing pro-diversity claims, moreover, reflected mainstream divisions in American society regarding the appropriateness of efforts to celebrate demographic diversity. People on the political left are more inclined than those on the political right to believe that demographic diversity should be celebrated (Fingerhut 2018). This partisan division was manifest in the contrasting ways in which Democratic and Republican politicians responded to the events. While Virginia’s Democratic governor Terry McAuliffe immediately condemned the rally, Republican president Donald Trump equivocated before famously affirming that “there were very fine people on both sides” (McAuliffe 2019). The emergence of diversity as a divisive sociopolitical issue is further evident in efforts under the Trump administration and Republican-led state governments to eliminate diversity, equity, and inclusion (DEI) initiatives (Fuchs, 2020), as well as in heated debates in local school board meetings regarding the appropriateness of educational initiatives related to racism, LGBTQ+ equality, and antisemitism (Thompson and Press 2022). In this sense, my approach mirrors recent research that has leveraged the January 6, 2021, assault on the U.S. capitol (Li and Disalvo 2023) or mass shootings (Hou and Poliquin 2022) to examine firms’ sociopolitical positioning. In these cases, as in the Unite the Right rally, sharply divided mainstream public opinion on the appropriate

response to fringe extremists put firms in a difficult position regarding whether they should speak up on the issue in question.

Comments from numerous interviewees suggested that Charlottesville's employers were aware of this tension as they transitioned from staying silent to taking pro-diversity positions. In the words of one employee,

Prior to that event, [my employer] had obviously had an advocacy group, which took positions on matters of public policy, but those were limited to its sort of sphere of, stay in our lane. We're going to talk about regulation. We're going to talk about market ethics. . . . We're not going to get out into broader social issues. And that changed for the organization as a result [of the rally], and there was some very robust internal discussion about where we should take positions on issues of social import.

Detailing this "robust internal discussion," this interviewee went on to describe spirited back and forth about whether the firm should embrace a pro-diversity stance, remain silent, or even speak up in favor of the protesters' First Amendment right to free speech. Similar debate was described by an interviewee from a different employer who indicated that it took an entire year of internal debate, paired with growing recognition that the rally was creating difficulties for recruiting, before the employer ultimately adopted pro-diversity claims.

Additional interviews and media reports provided further examples of employers across Charlottesville responding to the rally by eschewing silence and taking pro-diversity positions. One interviewee explained,

I think one of the things that really came out of the events from Unite the Right is that organizations I think went from a more passive support of social positions and sort of a tacit, we all got this right, to a much more active and vocal position on issues.

In an interview with the local newspaper *C'Ville Weekly*, Andrea Copeland, of the Charlottesville Chamber of Commerce, characterized this shift from silence to pro-diversity positioning as employers' effort to combat presumptions that they shared the attitudes that had been on display at the rally. Her words closely mirror my theoretical argument:

A lot of companies, for good reason, have sat down and said, "What can we do to make things better, to make everyone feel welcome, to make this community aware that we don't stand for exclusion, we don't stand for racism, we don't stand for what was on display August 11th and 12th?" (Pullinger 2019)

An interviewee at a local financial services company reported that whereas speaking up on social issues "had not been [the employer's] style in the first four years [the interviewee] worked there," after the rally that changed as the employer "came out with statements [about diversity] and put up some kind of fluffy, but nice stuff on our website." In another case, Jim Hall, CEO of WorldStrides, which sells educational trips for student groups, penned a letter titled "Our Hometown" on the company's website, in which he stated,

It was a disturbing and sad weekend here in WorldStrides' hometown of Charlottesville...Today especially, I find solace in the work that we do as an organization. We expose young people to new places, new ideas, new people, and new cultures. Through that, I believe that we help them develop an understanding, appreciation, and love of diversity, and the richness that it brings to our lives.

Teachstone, a company that develops and markets online training resources, expressed similar sentiments. In a letter titled "A Message about our Home" and signed "The Teachstone Family," the company affirmed the importance of "developing true and real relationships among people of all races, ethnicities, religions, and backgrounds." In another case, the Charlottesville-based Center



for Non-Profit Excellence lamented in an open letter on its website that “we were invaded this weekend. [sic] By violent, extremist hate groups that don’t reflect our collective values or the way we view the world.” The University of Virginia (UVA), Charlottesville’s largest employer, responded to the events with a series of countervailing pro-diversity claims. A general communication from the university stated that “Those who gather with the intent to strike fear and sow division do not reflect the University’s values and will not influence or diminish the University’s commitment to inclusion, diversity and mutual respect.” Shortly thereafter, a similar statement was issued by then-UVA president Teresa Sullivan, who affirmed that “the ideologies and beliefs expressed by many of the groups that have converged on Charlottesville this weekend contradict our values of diversity, inclusion, and mutual respect.” Illuminating a key motivation behind these statements, an interviewee who works in recruiting at UVA explained that in the wake of the rally, “we wanted to do everything we could to dispel the myth that this is who we are.” These many examples are consistent with my proposed theory: the rally created an anti-diversity stigma by association, which these employers felt compelled to combat via countervailing pro-diversity claims.

## **2.4 Econometric Analysis**

These qualitative anecdotes are consistent with my theory, but they do not allow me to systematically demonstrate the proposed effect or address alternative explanations. In this section, I use a difference-in-differences design on a comprehensive dataset of online job postings in the United States to estimate the rally’s effect on Charlottesville-area employers’ propensity to make pro-diversity claims in their recruiting efforts.

### **2.4.1 Countervailing Pro-Diversity Claims**

*Pro-diversity claims in job postings.* I operationalized sociopolitical claims in terms of firms' use of pro-diversity claims in job postings. Whereas White supremacy celebrates racial and cultural uniformity, pro-diversity claims carry a multicultural ideology that celebrates racial and cultural differences (Plaut et al. 2011, Dover et al. 2016). In 2017 in the United States, pro-diversity claims marked (and continue to mark) a sociopolitical stance in opposition not only to a historical, White-dominated status quo but to political forces, especially on the right, that resist trends in growing diversity (Dover et al. 2016, McVeigh and Estep 2019, Norris and Inglehart 2019). An analysis of corporate sociopolitical stance-taking from this period concluded that “diversity [was] the most frequently advocated issue” (Larcker et al. 2018: 2), while experimental research demonstrated that pro-diversity claims in job postings carry a strongly Democratic partisan connotation (Hurst 2021). Illustrating the sociopolitical nature of pro-diversity claims, a CNN photographer at the Unite the Right rally captured a now-iconic image of Confederate flag-bearing White supremacists marching far below a massive banner, hung by Charlottesville residents, emblazoned with the words “Diversity makes us stronger” (see Figure A1 in the Online Appendix). Thus, although some past research has characterized pro-diversity claims simply as pro-social claims (McDonnell and King 2013: 397), in this context they represent a position on one side of a divisive social issue.

Examining pro-diversity claims in the context of online job postings has two key advantages for evaluating my theorized relationship. First, job postings constitute a near-universal medium whereby firms can make sociopolitical claims. Whereas only some organizations have a meaningful media presence, nearly all organizations post jobs online. Second, geographic variation in job postings allows for fine-grained inference regarding the strategic behavior I theorize. In theory, firms can tailor job postings across time, locations, and labor pools in response to variation in sociopolitical stigma and worker preferences. Other mediums firms use to make sociopolitical

claims, such as social media posts or political donations, tend to lack location and labor pool specificity, making it more difficult to isolate my theorized mechanism.

To measure firms' use of pro-diversity claims, I used a dataset of all online job postings in the United States in the four-year period straddling the rally. I purchased these data from Burning Glass Technologies, which collects all job postings from approximately 45,000 company websites, online newspapers, and online job boards. The company ensures that no more than 5 percent of vacancies found this way come from any one source, and it removes duplicate postings across job boards. Third-party analysis from 2014 estimated that Burning Glass captures all job postings in the United States that appear online (Carnevale et al. 2014). Analysis by Burning Glass from 2022 indicates that dating back to 2013, the data capture 98.7 percent of jobs included in the U.S. Bureau of Labor Statistics Job Openings and Labor Turnover Survey (Lightcast 2022), which is the primary measure of U.S. job openings and is used extensively by policymakers and other labor market analysts. The data suffer some under-representation for jobs for which recruiting is less likely to occur online, such as those related to food preparation and service, office and administrative support, and construction. Thus, I can measure changes in employers' use of pro-diversity claims only in online recruiting.

I coded a job posting as featuring a pro-diversity claim if it met any one of these criteria: (1) it included the word "diversity" but not only within the words "diversityjobboard," "diversityjobs," "diversitynursing," or "pharmadiversity", (2) it included the word "inclusion" but not only when followed by "in," "of," "onto," or "with", (3) or it included the word "inclusive" but not only in the phrase "all inclusive" or "inclusive of." I detail, in the Online Appendix, the iterative process whereby I arrived at and validated this measure. This strategy is based on extensive analysis of random samples of job postings including the words "diversity," "equity,"

“inclusion,” “diverse,” “equitable,” or “inclusive.” I also examined a range of other words and phrases that might signal pro-diversity claims, including “belonging,” “bipoc,” “black lives matter,” “dei,” and “social justice,” but I found that these were rare. This approach yielded not only a very high true positive rate (I estimate that 96 percent of postings coded as featuring pro-diversity claims actually feature such claims) but also a very high true negative rate (I estimate that 97 percent of job postings coded as not featuring pro-diversity claims do not feature such claims).

*Additional variables.* I used additional variables at the job posting, employer, and geographic levels. At the job posting level, these include whether the job posts a salary (rather than an hourly wage or no information on compensation), the two-digit NAICS industry code, whether the job requires a bachelor’s degree, and the county in which the job is located. At the employer level, I identified employers that were not only hiring in Charlottesville but also based in or headquartered in Charlottesville. At the geographic level, I used the U.S. Census Bureau’s American Community Survey to collect annual five-year estimates of the county’s median age, the percentage of residents who identify as non-Hispanic White, the population, the fraction of residents who hold a bachelor’s degree, residents’ average annual earnings, and the vote share that went to Hillary Clinton in the 2016 election. I omitted job postings that were missing values on any of these job- or county-level variables or for which no employer was listed.

#### 2.4.2 Appropriate Level of Analysis

When making pro-diversity claims, employers may either tailor the language of individual postings or make one-off decisions whereby they standardize language across many or all postings. This variation raises the question of whether analyzing individual postings or individual employers is the appropriate strategy for capturing employers’ decisions to use pro-diversity claims. I feature

posting-level analysis as my main results for two related reasons. First, empirical patterns suggest that universal, standardized use of pro-diversity claims is the exception rather than the rule. Appendix Figure A2 shows that for the same employer, in the same county, in the same year, the share of job postings that feature pro-diversity claims varies significantly. Specifically, among employer–county dyads that make a pro-diversity claim in any of their postings in a given year, the percent of employers featuring pro-diversity claims in all of their postings is always below 30 percent and is well below 25 percent in three of the four years in my sample. Insights gleaned from analyzing this within-employer, posting-level variation helped me to arbitrate between competing theoretical explanations (see the section “Examining the Proposed Mechanism” below).

Second, by analyzing individual postings I also captured the reality that especially within larger employers, recruiting was not the purview of a single, central decision maker but, rather, the particular department, team, or division that was hiring. For example, although nearly all job postings by the University of Virginia (UVA) were coded by Burning Glass under a single employer, hiring was the purview of disparate academic departments, athletic teams, and administrative divisions. This likely explains, in part, why some but not all UVA postings across any six-month period in my data featured pro-diversity claims (see appendix Figure A4). Treating employers as the unit of analysis would inaccurately characterize larger employers as governed by a single decision maker or set of decision makers, giving them equal weight to smaller employers where this decision-making structure exists, such as single-location nail salons, restaurants, and laundromats.

There are two drawbacks to using job postings as the units of analysis. First, posting-level analysis arguably puts too much weight on employers that decide to include pro-diversity claims in most or all postings and then post many jobs. Second, posting-level analysis involves comparing

different populations of firms in the pre- and post-treatment periods. I used employer–county fixed effects to ensure that postings by employers that do not appear in the pre- and post-treatment periods do not directly contribute to the estimated treatment effects, but postings by these employers might indirectly shape this estimate through their influence on time fixed effects or coefficients on controls. To assuage these two concerns, I supplemented my principal analysis, in which the outcome is whether a given job posting features a pro-diversity claim, by analyzing a balanced-by-year panel of employer–county dyads, for which the outcome variable is the share of job postings in a given employer–county–year that feature pro-diversity claims. Although neither level of analysis flawlessly captured a decision to make a pro-diversity claim, I show that my results are robust to either approach.

#### 2.4.3 Summary Statistics

Table 1 presents summary statistics at the posting, employer, and county levels. I present these for the full sample, the Charlottesville MSA, counties that are home to flagship state universities, and counties in which more than 50 percent of the vote share went to Hillary Clinton in the 2016 presidential election.<sup>1</sup> Ninety-one percent of job postings in the Charlottesville MSA are in the city of Charlottesville. My full sample includes 66.5 million job postings, 12 percent of which feature a pro-diversity claim. This low rate is consistent with the idea, central to my theory, that employers are sensitive to the downsides associated with sociopolitical claims and thus tend to avoid them. Equal Opportunity claims, in contrast, which signal intentions for regulatory compliance, appear in 40 percent of postings. Consistent with my theorized relationship, Figure 2 shows that the share of Charlottesville-area job postings that featured pro-diversity claims increased 401 percent, from

8.3 percentage points in the pre-treatment period to 41.6 percentage points in the second year following the rally, while the use rate in jobs outside of Charlottesville increased only slightly. Whereas in the six months prior to the rally, the Charlottesville MSA was in the 64th percentile of MSAs in terms of the share of job postings that featured pro-diversity claims, two years later it had risen to the 99th percentile (see appendix Figure A3). Table A2 in the Online Appendix reports summary statistics for the employer-level, balanced panel dataset, and Figure 3 shows a similar increase in the use of pro-diversity claims when employers, rather than postings, are the units of analysis.

#### 2.4.4 Difference-in-Differences Analysis

I estimated a difference-in-differences model in which I compared the difference in the use rate of pro-diversity claims among job postings in the Charlottesville MSA to the use rate of pro-diversity claims among all other job postings in the pre- versus post-shock periods. Specifically, I estimated,

$$Diversity_{i,c,t} = \beta Charlottesville_{i,c,t} \times PostRally_{i,c,t} + X_{i,c,t} + \lambda_{i,c} + \lambda_t + \epsilon_{i,c,t} \quad (1)$$

where  $Diversity_{i,c,t}$  is a dummy equal to one if a job posting by employer  $i$  in county  $c$  in year-month  $t$  included a pro-diversity claim,  $Charlottesville \times PostRally$  is a dummy equal to one if the posting was made in the Charlottesville MSA after August 12, 2017,  $X$  is a vector of the job- and county-level controls listed above,  $\lambda_{i,c}$  are employer-county fixed effects,  $\lambda_t$  are month-year fixed effects where months are measured in relation to the rally (i.e., August 12, 2017), and  $\epsilon$  is the error term.  $\beta$  is the coefficient of interest and represents the difference-in-differences estimate of the Unite the Right rally's effect on the extent to which job postings in the Charlottesville MSA featured pro-diversity claims. I estimated this relationship using a linear probability model with robust standard errors clustered on county. Results, presented in Model 1 of Table 2, illustrate that the estimated treatment effect of the Unite the Right rally on the share of Charlottesville-area job

postings featuring pro-diversity claims was 8.6 percentage points ( $p < 0.001$ ). Relative to the unconditional pre-treatment baseline rate of 8.3 percentage points, this represents a sizeable increase of 104 percent.

#### 2.4.5 Evaluating Plausibility of the Parallel Trends Assumption

A key identifying assumption for causal interpretation of this estimate is the parallel trends assumption: in the absence of the Unite the Right rally, the difference between the rate at which job postings in the Charlottesville area featured pro-diversity claims and the rate at which postings in all other areas featured pro-diversity claims would have remained constant from the pre- to post-treatment periods. I conducted three exercises to examine the plausibility of this assumption.

**Placebo thresholds.** First, I re-estimated the treatment effect at the pre-treatment, placebo dates. I limited the sample to all job postings made in the two years prior to the rally and then re-estimated Equation 1, replacing the true treatment threshold with placebo treatment thresholds at three pre-treatment dates six months apart (February 11, 2016, August 11, 2016, February 11, 2017). These estimates, along with the estimate of  $\beta$  at the true treatment threshold for the full sample, are charted in appendix Figure A5. Although these three coefficients are significantly different from zero, they are all much smaller in magnitude than my estimated effect and in two of the three cases are in the opposite direction.

**Time-specific effects.** For a second test, I estimated an event-study model to examine whether there was a non-zero trend in the difference-in-differences estimates prior to treatment. Specifically, I estimated,

$$Diversity_{i,c,t} = \sum_{k \geq -3, k \neq 0}^{k=4} \delta_k Cville_{i,c,t} \times period_k + X_{i,c,t} + \lambda_{i,c} + \lambda_t + \epsilon_{i,c,t} \quad (2)$$

where *period* represents year dummies equal to one if the posting was made in period  $k$  where  $k \in [-3, 4]$  corresponds to each of the eight six-month periods in the study's time window (meaning  $k$



= -3 for entries made between August 12, 2015 and February 11, 2016, and  $k = 4$  for entries made between February 12, 2018 and August 11, 2019). The parameters  $\delta_k$  represent the period-specific difference-in-differences estimates. Note that the dummy for  $k = 0$  is omitted so that these estimates are relative to the six-month period before the rally (February 13, 2017 to August 12, 2017). If there were no confounding pre-trends, then the test should not reject the null that  $\delta_k = 0$  for all  $k < 0$ . Figure 4 shows these estimates and shows that two of these coefficients are significantly different from zero but are very small in magnitude. I also estimated this event-style model while excluding postings from UVA. Appendix Figure A6 presents these results. In this case, the effects are smaller but also climb in the post-treatment period.

***Parallel trends sensitivity analyses.*** Given the evidence of possible non-zero differences in pre-trends shown in Figure 4, I conducted a sensitivity analysis as suggested in Roth (2021). Roth recommended estimating the robustness of results to the hypothetical non-zero difference in pre-trends that would be detected 50 and 80 percent of the time. Appendix Figures A7 and A8 plot these hypothesized pre-trends relative to the estimates of  $\delta_k$  and show that, even allowing for non-zero pre-trends of this size, the post-treatment coefficients  $\hat{\delta}_2 - \hat{\delta}_4$  remain significant. Figures A9 and A10 show that the results excluding UVA (shown in Figure A6) are also robust to these hypothetical non-zero pre-trends. Together, these three analyses are consistent with the conclusion that the estimated effects cannot be attributed to non-parallel pre-trends.

#### 2.4.6 Considering Alternative Control Geographies

In the preceding analysis, the control group includes any job posting made outside of the Charlottesville MSA. A confounding trend or event, however, simultaneous with but separate from the Unite the Right rally, could have systematically increased the use of pro-diversity claims in

job postings in a subset of places similar to Charlottesville. To address this possibility, I identified alternative control geographies that, in different ways, are particularly similar to Charlottesville.

***Flagship state university counties.*** First, I identified counties that, like Charlottesville, are home to flagship state universities. Table 1 shows that these areas are more similar than the full U.S. sample to Charlottesville on numerous observed covariates. They may also be similar in important, less observable ways, such as their cultural importance and sensitivity to diversity issues. I re-estimated Equation 1, limiting the sample to postings in this subsample of locations. Model 2 of Table 2 presents the results. The estimated treatment effect in this case ( $\hat{\beta} = 0.086, p < 0.001$ ) is very similar to the estimate from the main specification.

***Clinton counties.*** Second, I identified counties in which, like Charlottesville, a majority of voters supported Hillary Clinton in the 2016 election. This addresses the possibility of a general pro-diversity backlash to Donald Trump's election in which employers in left-leaning areas augmented their use of pro-diversity claims. I again re-estimated Equation 1, limiting the sample to postings in this subsample of geographies. Model 3 of Table 2 presents the results of this regression. The estimated treatment effect in this case ( $\hat{\beta} = 0.089, p < 0.001$ ) is again very similar to the estimate from the main specification.

***UVA versus other higher education institutions.*** Third, I re-estimated Equation 1 and excluded all Charlottesville-area job postings not made by UVA and all job postings outside of Charlottesville that were not made by colleges, universities, and professional schools. This addresses the possibility of an exceptional upward trend in the use of pro-diversity claims among higher education institutions during my time period of analysis, which, since UVA features so prominently in the Charlottesville labor market, might confound my results. Model 4 of Table 2 reports these results. The effect in this case is large ( $\hat{\beta} = 0.215, p < 0.001$ ), representing a 140

percent increase from UVA’s pre-treatment use rate of pro-diversity claims. This result demonstrates that UVA was exceptional relative to other higher education institutions.

***Excluding UVA.*** Fourth, I again excluded UVA job postings and then re-estimated Equation 1. These results are reported in Model 5 of Table 2, which shows that the effect shrinks ( $\hat{\beta} = 0.024$ ) but is still highly significant ( $p < 0.001$ ). Relative to the unconditional pre-treatment baseline rate of pro-diversity claims among postings by non-UVA Charlottesville employers (0.066), this result still represents a sizeable increase, at 36 percent, in the use of pro-diversity claims. This estimate, however, indicates that UVA has a large effect on my posting-level estimates. For completeness, to give equal weight to all employers, I conducted the firm-level analysis detailed below in the “Employer-Level Analysis” section.

***Synthetic control comparison.*** Finally, I employed the synthetic control method proposed in Abadie et al. (2010), which involves comparing a treated unit, in this case Charlottesville, with a synthetic control unit created from a weighted average of the untreated units. In addition to providing a way to identify the control group, this method allowed me to address the possibility that because I have few treated units and many control units, errors are heteroskedastic. To execute this approach, I limited the comparison set to job postings in MSAs, collapsed the data to the MSA–month level, and then calculated the weights for the synthetic control, using the geographic-level controls from Equation 1. Appendix Table A7 demonstrates the improvement in balance achieved by this synthetic control method. Figures A11 and A12 illustrate the main result, showing that the difference in the use rate of pro-diversity claims in Charlottesville versus the synthetic Charlottesville rose in the years following the Unite the Right rally. I then conducted the randomization inference method developed by Abadie et al.(2010). I iteratively applied the synthetic control method to each untreated MSA. A sample of these placebo effects, along with

the effects for Charlottesville, is plotted in Figure A13. I then calculated the average mean squared prediction error (MSPE) for each placebo geography across the pre-treatment period, did the same across the post-treatment period, and calculated the percentage of placebo post- over pre-treatment MSPEs ratios that are greater than this ratio for Charlottesville. Figure A14 shows the distribution of this test statistic, which corresponds to a  $p$ -value of 0.010.

#### 2.4.7 Employer-Level Analysis

So far, I have treated individual postings as the units of analysis. As I explained above, this approach better captures decisions to make pro-diversity claims but has the potential disadvantages of over-weighting the decisions of larger employers and comparing different employers in the pre- and post-treatment periods. Accordingly, I repeated my main analysis at the employer level, using a balanced-by-year panel of employer–county dyads. Stated differently, the panel includes employer–county–year combinations in which the employer posted at least one job in that county in each of the four years of my study’s period. Here, the outcome of interest is the share of job postings in a given year in which a given employer in a given county made a pro-diversity claim. I estimated a two-way fixed effects difference-in-differences model,

$$DiversityShare_{i,c,y} = \beta Charlottesville_{i,c,y} \times PostRally_{i,c,y} + \lambda_{i,c} + \lambda_y + \epsilon_{i,c,y} \quad (3)$$

where  $DiversityShare_{i,c,y}$  is the share of job postings by employer  $c$  in county  $i$  in year  $y$  that featured a pro-diversity claim. The parameter of interest,  $\beta$ , represents the effect of the Unite the Right rally on the amount by which Charlottesville employers, on average, changed the share of their job postings that included pro-diversity claims. Results are presented in appendix Table A3. Here, the effect size is smaller ( $\hat{\beta} = 0.039$ ,  $p < 0.05$ ) but, relative to the pre-treatment unconditional baseline use rate of 9.3 percent, represents a sizeable increase of 42 percent. Table A4 shows the effects by year, again showing that the effect grew with time. The estimated effect for the final

year ( $\hat{\beta} = 0.053, p < 0.01$ ) represents a 58 percent increase from the unconditional pre-treatment baseline mean. For robustness, I again re-estimated the effect but limited the control group to counties that feature flagship state universities, limited the control group to counties in which over half of voters supported Hillary Clinton in the 2016 election, limited the comparison to UVA versus other higher education institutions, and then removed UVA postings from the sample. Models 2–5 of Table A4 show that the results are robust in these alternative specifications.

#### 2.4.8 Stable Unit Treatment Value Assumption Violations

Another identifying assumption in this context is the stable unit treatment value assumption (SUTVA): for my estimates to be unbiased, the rally would have affected Charlottesville exclusively and not spilled into other geographies. SUTVA is likely violated inasmuch as the rally received extensive media attention, plausibly resulting in a broader generalization process wherein job seekers became more inclined to presume anti-diversity stances among employers in locations in addition to Charlottesville. The bias resulting from this violation, however, would likely attenuate the estimated treatment effects inasmuch as, on balance, employers in these other locations, especially those in my alternative control subsets, would similarly tend to rely on left-leaning workers and would thus similarly face incentives to use countervailing pro-diversity claims to combat any spillover presumptions of anti-diversity positions.

### **2.5 Examining the Proposed Mechanism**

I have proposed that the sizeable increase in employers' use of pro-diversity claims is a strategic, defensive response by Charlottesville employers meant to countervail the anti-diversity stigma they incurred from their proximity to the Unite the Right rally. But outside of the qualitative anecdotes described above, I cannot directly observe whether this was employers' motivation for

adopting pro-diversity claims. Moreover, my theorized mechanism is unlikely the only motivation for this behavior. In this section, I present a range of empirical patterns that are consistent with my proposed mechanism and less consistent with alternative explanations, including isomorphic processes, increased salience of diversity issues, upper-echelon attitudes, or efforts to comply with the Equal Opportunity Act. Table 3 summarizes these alternative explanations and empirical patterns that are more consistent with my theorized mechanism than with these other explanations.

### 2.5.1 Examining Effects by Headquarters Location

If employers' adoption of pro-diversity claims was strategic rather than the result of a more general isomorphic process or simply a manifestation of top managers' values, it should have been greatest among employers most at risk of incurring anti-diversity stigma by association. To test this, I examined whether adoption was larger for employers not only located in Charlottesville but also headquartered there. This differentiates, for example, UVA and the CFA Institute, which are based in Charlottesville, from Bank of America, which has locations in Charlottesville but is headquartered in Charlotte, North Carolina. Since Charlottesville-based employers are more socially proximate to Charlottesville, they plausibly face greater risk of incurring stigma by association and, accordingly, face an outside incentive to combat this stigma through countervailing pro-diversity claims. I estimated,

$$\begin{aligned}
 Diversity_{i,c,t} = & \beta_1 CVilleBased_{i,c,y} \times Charlottesville_{i,c,t} \times PostRally_{i,c,t} + \\
 & \beta_2 CVilleBased_{i,c,t} + \beta_3 CVilleBased_{i,c,t} \times Charlottesville_{i,c,t} + X_{i,c,t} + \lambda_{i,c} + \lambda_t + \epsilon_{i,c,t}
 \end{aligned}
 \tag{4}$$

where *CVilleBased* is a dummy equal to one if the job posting was made by an employer based in Charlottesville. In this case,  $\beta_1$  is the coefficient of interest. These results are presented in Model 1 of Table 4, which shows that the effect of the Unite the Right rally on the use of pro-diversity

claims in job postings was significantly larger for Charlottesville-based employers ( $\hat{\beta}_1 = 0.176, p < 0.001$ ). To examine whether this effect may have been driven entirely by employers such as UVA that posted many jobs, I repeated this analysis using the balanced panel, estimating a version of Model 3 in which I interacted the treatment dummy with *CvilleBased*. These results, reported in Model 1 of Table A5, again show that the effect is significantly larger for employers headquartered in Charlottesville ( $\hat{\beta}_1 = 0.056, p < 0.01$ ).

### 2.5.2 Examining Effects by Target Applicant Pro-Diversity Preferences

I also examined whether the increase in pro-diversity claims in job postings was largest when prospective applicants held the strongest pro-diversity preferences. I tested this by leveraging the fact that bachelor's-holding job seekers, on average, held significantly stronger pro-diversity views than did their non-bachelor's-holding counterparts (Parker et al. 2019). This education-based division was starkly visible in the 2016 presidential election, in which non-college graduates went 50/43 percent for Donald Trump, whose campaign became synonymous with anti-diversity positions and who was famously reticent to condemn the Unite the Right rally, while college graduates went 57/36 percent for Clinton (Pew Research Center 2018).

To run this test, I re-estimated Equation 4 but replaced *CvilleBased* with a dummy equal to one if the posting indicated that applicants must hold at least a bachelor's degree. Model 2 of Table 4 presents these results and shows that the effect was significantly larger ( $\hat{\beta}_1 = 0.028, p < 0.001$ ) for such jobs. I again repeated this at the employer level, using the balanced panel. These results are reported in Model 2 of Table A5 and show the effect was significantly larger ( $\hat{\beta}_1 = 0.046, p < 0.001$ ) among employers that were in the top quartile of employer–county dyads in terms of the share of their jobs that required a bachelor's degree, which corresponds to employers for which more than 39 percent of jobs required a bachelor's degree.

### 2.5.3 Examining Effects by Geographic Scope of Recruiting

An alternative strategic explanation for the increased use of pro-diversity claims is that the main effects simply reflected firms' efforts to appeal to Charlottesville's workers, for whom, due to their proximity to the rally, issues of diversity had suddenly become very salient. If this were the case, the effects should be relatively larger when the geographic scope of recruiting was mostly limited to Charlottesville and smaller when employers sought applicants outside of Charlottesville. My proposed explanation, in contrast, would suggest that these effects would be largest when recruitment efforts focused on prospective employees *outside* of Charlottesville because those employees were less aware that the rally's agitators were mostly from outside of Charlottesville, and thus those employees were more susceptible to my theorized generalization process. Multiple Charlottesville-based interviewees substantiated this variation in perceptions, expressing frustration that due to the rally, outsiders had inaccurately associated the community with anti-diversity positions.

To examine this possibility, I estimated heterogeneous treatment effects with respect to whether the job posting included a commitment to cover moving expenses. It is highly plausible that postings featuring these commitments were more likely to seek applicants outside of Charlottesville, compared to those that did not. I re-estimated Equation 4 but replaced *CVilleBased* with a dummy equal to one if the job posting included a commitment to cover moving expenses.

Model 3 of Table 4 presents the results and shows that the effect in these cases is significantly larger ( $\hat{\beta}_1 = 0.165, p < 0.001$ ). Again, I repeated this analysis at the firm level, using the balanced panel. Model 3 of Table A5 shows that the effect was larger ( $\hat{\beta}_1 = 0.242, p < 0.10$ ) among employers for whom over half of their job postings included a commitment to cover moving expenses. Besides addressing this alternative salience explanation, this test is further evidence of a



strategic rather than isomorphic or upper-echelon perspective in that the adoption of pro-diversity claims was again larger where employers faced greater incentives to make such claims.

#### 2.5.4 Examining Intrafirm Geographic Variation

Next, I conducted another test to arbitrate between my theorized strategic explanation and an upper-echelon explanation that would characterize these patterns as an expression of top managers' pro-diversity beliefs. While holding top managers constant, I examined whether a given employer was more likely to adopt pro-diversity claims in its Charlottesville job postings compared to its postings in other locations. I did so by estimating a version of Equation 1 in which I included company-year fixed effects (in addition to month and company-county fixed effects). This means that postings from employers that posted only in one county do not contribute directly to the estimated effect, which is why I did not feature this in my main analysis. I report these estimates in Model 1 of Table 5, which shows that in this case the effect is smaller but still significant ( $\hat{\beta} = 0.010, p < 0.01$ ). This finding is consistent with the idea that employers selectively adjusted their recruiting language to countervail the anti-diversity stigma that had arisen uniquely in Charlottesville. As Table A6 reports, I repeated these analyses at the employer level and find a relationship that is of similar size ( $\hat{\beta}_1 = 0.011, p = 0.16$ ) but not statistically significant.

#### 2.5.5 The Delayed Effect

As shown in the posting-level and employer-level analysis in both descriptive statistics (see Figures 2 and 3) and in difference-in-difference estimates (see Figures 4 and A6 as well as Table A4), the increased use of pro-diversity claims did not arise immediately following the rally but began to materialize after approximately six months and grew throughout the post-treatment period. If the effect had been driven primarily by an increase in the salience of diversity issues

within Charlottesville or by top managers' newfound desire to express their pro-diversity values, it would have been more likely to materialize immediately and then stabilize or diminish as the issue decreased in salience and/or top managers' attention was drawn to other issues. The delayed effect, in contrast, is more consistent with the idea that the increased use of pro-diversity claims represented a calculated response from employers. First, the delay is consistent with a learning process wherein employers did not immediately appreciate the difficulties the rally had created for their recruiting efforts but, as prospective employees consistently mentioned the rally, began making pro-diversity claims. This learning process was substantiated by multiple interviewees who described their initial surprise when, following the rally, prospective employees began inquiring about employers' stances on diversity. Second, the delay is consistent with the idea that employers deliberated and debated prior to their decisions to adopt pro-diversity stances. This, too, was substantiated by multiple interviewees.

#### 2.5.6 Significant Non-Adoption

While the increased use of pro-diversity claims represented a substantial shift relative to the pre-treatment use rate in Charlottesville, by the end of the study period, a majority of postings in Charlottesville still did *not* feature pro-diversity claims. This pattern provides further support for my theorized mechanism. Universal or near-universal adoption of pro-diversity claims would be more consistent with an isomorphic process wherein the Unite the Right rally caused a general diffusion of the practice or initiated a general norm among Charlottesville employers to use these claims. Instead, I demonstrate that the adoption of the practice is selective and concentrated among employers that had the greatest incentive to do so, including those that were more likely to be associated with the Unite the Right rally and were targeting job seekers who held strong pro-diversity preferences.

### 2.5.7 Examining Equal Opportunity Claims

A separate alternative explanation is that this increase in pro-diversity claims mostly resulted from efforts to signal compliance with the Equal Opportunity Act (EOA), which prohibits employers from discriminating based on race, color, sex, religion, or national origin (Dobbin, 2009). Employers may have feared that the Unite the Right rally made employees, job seekers, regulators, or even the private bar more inclined to seek legal action for perceived discrimination, a possibility consistent with research illustrating that social activism often shapes broader regulatory outcomes (Fremeth et al. 2022). Since Equal Opportunity (EO) claims are often accompanied by pro-diversity claims, an increase in EO claims motivated by a desire to signal compliance with the EOA may have mechanically led to an increase in pro-diversity claims. To examine this possibility, I re-estimated Equation 1 but replaced the outcome *Diversity* with a dummy equal to one if the job posting included an EO claim. I coded a job posting as including an EO claim if it featured the phrase “equal opportunity,” “equal employment opportunity,” or “eo” as a standalone abbreviation. Model 2 of Table 5 shows a significant increase in EO claims but that this effect was much smaller ( $\hat{\beta}=0.035$ ) than the overall increase in pro-diversity claims ( $\hat{\beta}=0.086$ ). Model 2 of Table A6 reports this same analysis at the employer level and shows an effect that is statistically insignificant and even smaller relative to the employer-level estimate of the increase in pro-diversity claims.

### 2.5.8 Mentioning Charlottesville in Job Postings

Finally, I examined changes in the extent to which job postings featured the word “Charlottesville.” Although this pattern is somewhat distinct from my theory of countervailing claim-making, a decrease in mentions would be consistent with my broader argument that firms recognized the liability created by their association with Charlottesville and sought to downplay

this association in their recruiting efforts. Consistent with this idea, Model 3 in Table 5 shows that there was a significant decrease ( $\hat{\beta} = -0.018, p < 0.001$ ) in the extent to which Charlottesville-area employers included the word “Charlottesville” in their job postings. Model 3 of Table A6 shows this effect is even larger ( $\hat{\beta} = -0.042, p < 0.001$ ) when analyzed at the employer level.

## 2.6 Supplementary Wage Analysis

To conduct a supplementary exercise, I examined whether the rally may have made it more difficult for Charlottesville’s employers to attract workers and whether pro-diversity claims ameliorated these difficulties. To examine whether the rally made it more difficult to attract workers, I focused on whether it affected the wages offered by Charlottesville’s employers. A newfound wage premium would be consistent with the idea that employers used higher wages to compensate for job seekers’ diminished inclination to work for them due to their stigma by association with the rally. To estimate the wage premium, I re-estimated Equation 1 but replaced the outcome *Diversity* with the annual salary listed in the job posting. When the annual salary was expressed as a range, I used the lower bound of this range. These results are presented in Model 1 of Table 6 and show that the rally is associated with a wage premium of \$2,500.

I next examined whether this premium was lower when employers used pro-diversity claims. A lower premium in the presence of pro-diversity claims would be consistent with the idea that these claims assuaged prospective employees’ concerns about the employers’ stances on diversity, allowing employers to substitute them for the higher wages they would otherwise pay to compensate for these concerns. In estimating this relationship, I acknowledge that since pro-diversity claims were not randomized, associations might have been caused by unaccounted-for third factors. The continued inclusion of a rich set of fixed effects and controls mitigates but does not eliminate this risk. I also cannot observe the relationship between different posting

characteristics (i.e., salary offered, diversity claims) and whether the employer successfully filled the advertised position. Understanding these limitations, I re-estimated a version of Equation 4 in which I replaced the outcome variable *Diversity* with the annual salary and replaced the interacted variable *CVilleBased* with *Diversity*. These results are presented in Model 3 of Table 6 and show that the use of pro-diversity claims is associated with a significant reduction of the wage premium. Table A8 shows that these patterns are robust to the exclusion of postings by UVA.

## **2.7 Discussion and Conclusion**

This article provides a novel strategic explanation for why firms take positions on divisive sociopolitical issues. Drawing on theories of stigma by association, I explained why firms' proximity to controversial political actors may lead stakeholders to presume that firms silent on sociopolitical issues are misaligned with the stakeholders' sociopolitical preferences and why firms combat these presumptions by making countervailing sociopolitical claims. I substantiated this theory by showing that after the 2017 Unite the Right rally, Charlottesville's employers combated presumptions that they shared demonstrators' anti-diversity positions by adopting countervailing pro-diversity claims in their online job postings. Consistent with my theory, I showed that the adoption of pro-diversity claims was not exceptionally widespread, it occurred most often when employers targeted prospective applicants who were plausibly more inclined to generalize the White supremacists' anti-diversity positions and who held strong pro-diversity preferences. I showed, moreover, that the rally was associated with a newfound wage premium for Charlottesville-area job postings but that this premium was lower for job postings featuring pro-diversity claims.

### **2.7.1 Strategic Sociopolitical Positioning**

This study advances nascent literature on strategic sociopolitical positioning, or the idea that firms' sociopolitical positioning represents calculated appeals to stakeholders. It differs from related work suggesting that firms may use social claims to combat negative evaluations arising from their own actions (Elsbach 1994, Bansal and Clelland 2004, McDonnell and King 2013) or to differentiate themselves from their competitors (Mohliver et al. 2023). I demonstrate that certain firms strategically differentiate themselves from non-firm political actors with whom they are associated and whom key stakeholders view unfavorably. Whereas a growing body of both empirical (Burbano 2021, Hou and Poliquin 2022, Wowak et al. 2022) and theoretical (Melloni, et al. 2023, Mohliver et al. 2023) research has sought to identify when sociopolitical positioning benefits firms, to my knowledge this article provides the first empirical evidence of firms strategically claiming sociopolitical positions. Further elaboration of this strategic perspective requires additional empirical studies that similarly treat firm behavior, rather than stakeholder reactions, as the key outcome of interest. Future studies may seek to theorize how strategic motivations, such as those outlined here and those outlined by Mohliver et al. (2023) and Melloni et al. (2023), interact with expressive motivations, such as those identified by Hambrick and Wowak (2021). For example, top managers might generally prefer to express their political beliefs but carefully limit these expressions to instances likely to benefit the firm.

By examining sociopolitical claims in the context of job postings rather than in the public statements commonly examined in past research, this article also demonstrates how a firm's positioning can be subtle and even tailored to particular stakeholder audiences. Future studies might compare firms' job posting language with messages to consumers, shareholders, or other stakeholders. Evidence of a focal firm making contrasting claims that correlate with stakeholders' contrasting preferences would further enrich the strategic view of sociopolitical stance-taking. Such

research will benefit by examining novel data. Surveys by Hertel-Fernandez (2018), for example, suggested that in the United States, nearly 50 percent of managers attempt to mobilize their workers politically and 25 percent of workers have had a manager try to mobilize them. Analyzing the internal communications whereby these mobilization efforts occur would provide important insights into sociopolitical positioning. More generally, a fuller account of strategic positioning might clarify whether and/or when stakeholders' reactions to perceptions of sociopolitical alignment are fundamentally taste-based (i.e., due simply to ideological aversion) or statistical (i.e., due to the belief that the sociopolitical position signals additional firm qualities). Such an account might elucidate, for example, the extent to which firms need to back sociopolitical claims with evidence that they behave in a way consistent with these claims.

Future work may also examine how these practices evolve in the context of growing political polarization. On one hand, growing polarization may generate ever-stronger negative reactions to firms' sociopolitical stance-taking, which would seem to disincentivize such positioning. On the other hand, growing polarization may lead to greater sorting by political ideology, a possibility suggested by Bermiss and McDonald (2018). Sorting by political ideology would, it seems, incentivize sociopolitical positioning since there would be greater within-firm political homogeneity and, thus, less risk of alienating people on the other side of the position. The result of polarization may thus be an environment in which a shrinking number of firms with politically heterogeneous stakeholder bases deliberately maintain silence while a growing number of firms with politically homogeneous stakeholder bases increasingly speak up.

### 2.7.2 Human Capital Strategy

This article makes secondary contributions to research on non-pecuniary human capital strategy. Prior research has demonstrated how non-pecuniary firm characteristics can enhance potential

applicants' inclination to apply to jobs (Abraham and Burbano 2021), increase the amount of effort exerted by workers (Carpenter and Gong 2015, Burbano 2021), improve employee retention (Bode et al. 2015, Carnahan et al. 2017, Berniss and McDonald, 2018), and even reduce knowledge spillovers (Flammer and Kacperczyk 2019). To date, however, this literature has been largely agnostic on whether firms cultivate these characteristics in an attempt to reap these benefits. To my knowledge, this study provides the first evidence of firms strategically adjusting their human capital strategy in an attempt to align with workers' exogenously given, non-pecuniary preferences. Moreover, although I provide preliminary evidence that employers tailor their pro-diversity language and seem to do so in strategic ways, future research could do more to theorize and test the conditions under which employers standardize or tailor diversity-related recruitment strategies as well as the effects of these actions. Prior studies, for example, have illustrated how multi-location firms can benefit by delegating product development and advertising responsibilities to local offices that are better positioned to respond to the tastes of local markets. Are there analogous benefits to delegating recruiting strategies to local human resource functions that may be similarly well positioned to appreciate the political preferences in a given labor pool?

### 2.7.3 Diversity Strategy

By operationalizing sociopolitical positioning with employers' use of pro-diversity language, this article also relates to literature on workplace inequality and diversity in recruiting (Cobb 2016, Tolbert and Castilla 2017, Dobbin and Kalev 2021, Hurst et al. 2022). Much of this research has examined organizational practices relating to current employees, such as diversity training, formalization of promotion criteria, and grievance systems (Kalev et al. 2006, Dobbin 2009). Diversity claims, in contrast, target prospective employees and have potentially large impacts on who applies to and, ultimately, works at the firm. And while several lab and field experimental studies have



demonstrated that these claims may in fact enhance applicant pool diversity (Kang et al. 2016, Abraham and Burbano 2021, Flory et al. 2021), I provide, to my knowledge, the most comprehensive estimates to date of the extent to which firms actually use these claims. I estimate that in my study's time frame, 12 percent of U.S. job listings featured pro-diversity claims. Given the ease with which firms can include pro-diversity claims in their recruiting efforts as well as growing public and regulatory pressure for workplace diversity, why is this rate so low? Do employers eschew these claims to avoid accusations of hypocrisy? Do they fear they will be seen as too political? I also present findings consistent with the conclusion that employers perceived as hostile toward minorities may need to pay a wage premium, but that these perceptions and their accompanying wage costs can diminish when employers make pro-diversity claims. Using the data and strategy for measuring pro-diversity claims that I present here, future research can address these questions and more fully illustrate when and to what effect employers use pro-diversity claims in their recruiting efforts.

#### 2.7.4 Locational Stigma and Implications for Integrated Strategy

Finally, this study advances research on stigma by association. Whereas past research has demonstrated how stigmatizing generalization processes disseminate across organizations that share a parent organization (Jensen 2006, Piazza and Jourdan 2018), partnership (McDonnell, et al. 2021), organizational form (Yue et al. 2013), or industry (Jonsson et al. 2009, Piazza and Perretti 2015, Naumovska and Lavie 2021), this article considers how these processes may disseminate due to simply sharing a location with non-firm or even non-organizational actors. This possibility has significant implications for integrated strategy, or the idea that firms enhance the success of their market strategies by undertaking complementary non-market strategies (Holburn and Vanden Bergh 2014, Barber IV and Diestre 2019). When, for example, a firm opens offices or acquires a

firm in a new location, it exposes itself to possible stigma arising from positions taken by political actors in that location. This study suggests conditions under which these location-based market strategies should be paired with sociopolitical claims in order to combat stakeholders' perceptions of sociopolitical misalignment.

These integrated strategies are visible in the U.S. in the case of corporate relocations to right-leaning states. Firms including Toyota, Hewlett Packard, and Tesla have recently moved headquarters from left-leaning states, such as California, to right-leaning Southern states, such as Texas. These locations offer attractive tax advantages and a low cost of living. In entering these locations, however, firms risk acquiring sociopolitical stigma resulting from right-leaning positions established by, among other things, anti-trans bathroom bills and abortion bans. Underlining this tension, *Fortune* magazine reported that although these states provide “pro-enterprise climates,” they are “simultaneously rolling out [right-leaning] laws that are an automatic turnoff to many young, liberal tech workers and force companies into damage control mode” (Leonhardt 2021). Anecdotal evidence suggests that this “damage control” involves combating right-leaning stigma with countervailing left-leaning sociopolitical claims (Gelles 2021, Pardes 2021). Future research might use this context to test whether firms systematically employ countervailing sociopolitical claims to combat sociopolitical stigma arising from mergers, acquisitions, or relocations into right-leaning states.

### 2.7.5 A Note on Generalizability

The exceptional nature of the Unite the Right rally raises questions about this study's generalizability. Even though there is surprisingly widespread sympathy with many ideas core to White supremacy (UVA Center for Politics, 2017), such dramatic manifestations are rare. Given the frequency with which politicians, legislatures, courts, electorates, and protesters take positions

viewed unfavorably by large populations of stakeholders, however, the strategic behavior I theorize in this paper is plausibly ubiquitous. Numerous anecdotes suggest this is the case. In the wake of Texas's Senate Bill 8, which codified a strongly pro-life stance on abortion, many Texas-based tech companies made countervailing commitments to women's and reproductive rights (Pardes 2021). In the wake of Donald Trump's 2016 victory, which many saw as revealing the ubiquity of anti-diversity sentiment in the U.S. electorate, many U.S. firms made countervailing commitments to racial, gender, and LGBTQ+ diversity (*The Economist* 2017). These and numerous similar contexts, however, present serious impediments to disentangling my theorized explanation from plausible alternative explanations. I overcame these impediments by exploiting idiosyncrasies of the Unite the Right rally. Thus, I leveraged idiosyncrasies of a somewhat unusual empirical context to test a broadly generalizable theoretical mechanism, which is similar to recent work by Favaron et al. (2022). And while I consider the case of human capital and labor markets, my broader theoretical argument plausibly generalizes to stakeholders in product and other-factor markets who, past research has shown, are similarly disposed to allow their political preferences to shape their decisions to reward or spurn firms (McConnell et al. 2018, Chatterji and Toffel 2019). Examining these additional stakeholder groups, moreover, may illuminate additional conditions under which countervailing claims are more or less likely to arise and/or achieve their intended purpose.

## 2.8 Works Cited

- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association* 105, no. 490 (June 2010): 493–505. <https://doi.org/10.1198/jasa.2009.ap08746>.
- Abraham, Mabel, and Vanessa Burbano. "Congruence Between Leadership Gender and Organizational Claims Affects the Gender Composition of the Applicant Pool: Field Experimental Evidence." *Organization Science* Articles in Advance (March 17, 2021): 1–25. <https://doi.org/10.1287/orsc.2021.1442>.
- Adam Cobb, J. "How Firms Shape Income Inequality: Stakeholder Power, Executive Decision Making, and the Structuring of Employment Relationships." *Academy of Management Review* 41, no. 2 (April 2016): 324–48. <https://doi.org/10.5465/amr.2013.0451>.
- Adut, Ari. "A Theory of Scandal: Victorians, Homosexuality, and the Fall of Oscar Wilde." *American Journal of Sociology* 111, no. 1 (July 2005): 213–48. <https://doi.org/10.1086/428816>.
- Argenti, Paul A. "Collaborating with Activists: How Starbucks Works with NGOs." *California Management Review* 47, no. 1 (October 1, 2004): 91–116. <https://doi.org/10.2307/41166288>.
- Bakare, Lanre. "Drama Schools Accused of Hypocrisy over Anti-Racism Statements." *The Guardian*, June 9, 2020, sec. Stage. <https://www.theguardian.com/stage/2020/jun/06/drama-schools-accused-of-hypocrisy-over-anti-racism-statements>.
- Bansal, Pratima, and Iain Clelland. "Talking Trash: Legitimacy, Impression Management, and Unsystematic Risk in the Context of the Natural Environment." *The Academy of Management Journal* 47, no. 1 (2004): 93–103. <https://doi.org/10.2307/20159562>.
- Barber IV, Benjamin, and Luis Diestre. "Pushing for Speed or Scope? Pharmaceutical Lobbying and Food and Drug Administration Drug Review." *Strategic Management Journal* 40, no. 8 (2019): 1194–1218. <https://doi.org/10.1002/smj.3021>.
- Baron, David P., and Daniel Diermeier. "Strategic Activism and Nonmarket Strategy." *Journal of Economics & Management Strategy* 16, no. 3 (2007): 599–634. <https://doi.org/10.1111/j.1530-9134.2007.00152.x>.
- Baumeister, Roy F., Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D. Vohs. "Bad Is Stronger than Good." *Review of General Psychology* 5, no. 4 (December 2001): 323–70. <https://doi.org/10.1037/1089-2680.5.4.323>.

- Bermiss, Y. Sekou, and Rory McDonald. “Ideological Misfit? Political Affiliation and Employee Departure in the Private-Equity Industry.” *Academy of Management Journal*, February 27, 2018. <https://doi.org/10.5465/amj.2016.0817>.
- Bhagwat, Yashoda, Nooshin L. Warren, Joshua T. Beck, and George F. Watson. “Corporate Sociopolitical Activism and Firm Value.” *Journal of Marketing* 84, no. 5 (September 1, 2020): 1–21. <https://doi.org/10.1177/0022242920937000>.
- Bode, Christiane, Jasjit Singh, and Michelle Rogan. “Corporate Social Initiatives and Employee Retention.” *Organization Science* 26, no. 6 (December 2015): 1702–20. <https://doi.org/10.1287/orsc.2015.1006>.
- Bolton, Patrick, Tao Li, Enrichetta Ravina, and Howard Rosenthal. “Investor Ideology.” *Journal of Financial Economics* 137, no. 2 (August 1, 2020): 320–52. <https://doi.org/10.1016/j.jfineco.2020.03.004>.
- Briscoe, F., M. K. Chin, and D. C. Hambrick. “CEO Ideology as an Element of the Corporate Opportunity Structure for Social Activists.” *Academy of Management Journal* 57, no. 6 (December 1, 2014): 1786–1809. <https://doi.org/10.5465/amj.2013.0255>.
- Briscoe, Forrest, Abhinav Gupta, and Mark S. Anner. “Social Activism and Practice Diffusion: How Activist Tactics Affect Non-Targeted Organizations.” *Administrative Science Quarterly* 60, no. 2 (June 1, 2015): 300–332. <https://doi.org/10.1177/0001839215579235>.
- Burbano, Vanessa C. “The Demotivating Effects of Communicating a Social-Political Stance: Field Experimental Evidence from an Online Labor Market Platform.” *Management Science* 67, no. 2 (February 2021): 1004–25. <https://doi.org/10.1287/mnsc.2019.3562>.
- Carlos, W. Chad, and Ben W. Lewis. “Strategic Silence: Withholding Certification Status as a Hypocrisy Avoidance Tactic.” *Administrative Science Quarterly* 63, no. 1 (March 1, 2018): 130–69. <https://doi.org/10.1177/0001839217695089>.
- Carnahan, Seth, David Kryscynski, and Daniel Olson. “When Does Corporate Social Responsibility Reduce Employee Turnover? Evidence from Attorneys Before and After 9/11.” *Academy of Management Journal* 60, no. 5 (October 2017): 1932–62. <https://doi.org/10.5465/amj.2015.0032>.
- Carnevale, Anthony, Tamara Jayasundera, and Dmitri Repnikov. “Understanding Online Job Ads Data: A Technical Report.” Technical Report. Georgetown University, Center on Education and the Workforce, April 2014.
- Carpenter, Jeffrey, and Erick Gong. “Motivating Agents: How Much Does the Mission Matter?” *Journal of Labor Economics* 34, no. 1 (December 21, 2015): 211–36. <https://doi.org/10.1086/682345>.

- Castilla, Emilio J. “Accounting for the Gap: A Firm Study Manipulating Organizational Accountability and Transparency in Pay Decisions.” *Organization Science* 26, no. 2 (April 1, 2015): 311–33. <https://doi.org/10.1287/orsc.2014.0950>.
- Chatterji, Aaron K., and Michael W. Toffel. “Assessing the Impact of CEO Activism.” *Organization & Environment* 32, no. 2 (June 1, 2019): 159–85. <https://doi.org/10.1177/1086026619848144>.
- Conover, Pamela Johnston, and Stanley Feldman. “Projection and the Perception of Candidates’ Issue Positions.” *The Western Political Quarterly* 35, no. 2 (1982): 228–44. <https://doi.org/10.2307/448017>.
- Devers, Cynthia E., Todd Dewett, Yuri Mishina, and Carrie A. Belsito. “A General Theory of Organizational Stigma.” *Organization Science* 20, no. 1 (February 1, 2009): 154–71. <https://doi.org/10.1287/orsc.1080.0367>.
- Dewan, Torun, and Kenneth A. Shepsle. “Political Economy Models of Elections.” *Annual Review of Political Science* 14, no. 1 (2011): 311–30. <https://doi.org/10.1146/annurev.polisci.12.042507.094704>.
- Dobbin, F. *Inventing Equal Opportunity*. Princeton University Press, 2009. <https://books.google.com/books?id=GCAz25CZKoIC>.
- Dobbin, Frank, and Alexandra Kaley. “The Civil Rights Revolution at Work: What Went Wrong.” *Annual Review of Sociology* 47, no. 1 (2021): 281–303. <https://doi.org/10.1146/annurev-soc-090820-023615>.
- Documenting Hate: The New American Nazis*. PBS, 2018.
- Dover, Tessa L., Brenda Major, and Cheryl R. Kaiser. “Members of High-Status Groups Are Threatened by pro-Diversity Organizational Messages.” *Journal of Experimental Social Psychology* 62 (January 1, 2016): 58–67. <https://doi.org/10.1016/j.jesp.2015.10.006>.
- Durand, Rodolphe, and Jean-Philippe Vergne. “Asset Divestment as a Response to Media Attacks in Stigmatized Industries.” *Strategic Management Journal* 36, no. 8 (2015): 1205–23. <https://doi.org/10.1002/smj.2280>.
- Elan, Priya. “Munroe Bergdorf Accuses L’Oréal of Racial Hypocrisy.” *The Guardian*, June 2, 2020, sec. Fashion. <https://www.theguardian.com/fashion/2020/jun/02/munroe-bergdorf-accuses-loreal-of-racial-hypocrisy>.
- Elsbach, Kimberly D. “Managing Organizational Legitimacy in the California Cattle Industry: The Construction and Effectiveness of Verbal Accounts.” *Administrative Science Quarterly* 39, no. 1 (1994): 57–88. <https://doi.org/10.2307/2393494>.

- Favaron, Saverio Dave, Giada Di Stefano, and Rodolphe Durand. “Michelin Is Coming to Town: Organizational Responses to Status Shocks.” *Management Science*, January 6, 2022. <https://doi.org/10.1287/mnsc.2021.4210>.
- Fingerhut, Hannah. “Most Americans Express Positive Views of Country’s Growing Racial and Ethnic Diversity.” *Pew Research Center* (blog), 2018. <https://www.pewresearch.org/fact-tank/2018/06/14/most-americans-express-positive-views-of-countrys-growing-racial-and-ethnic-diversity/>.
- Flammer, Caroline, and Aleksandra Kacperczyk. “Corporate Social Responsibility as a Defense against Knowledge Spillovers: Evidence from the Inevitable Disclosure Doctrine.” *Strategic Management Journal* 40, no. 8 (2019): 1243–67. <https://doi.org/10.1002/smj.3025>.
- Flory, Jeffrey A., Andreas Leibbrandt, Christina Rott, and Olga Stoddard. “Increasing Workplace Diversity: Evidence from a Recruiting Experiment at a Fortune 500 Company.” *Journal of Human Resources* 57, no. 1 (2021): 73–92. <https://doi.org/10.3368/jhr.56.1.0518-9489R1>.
- Fremeth, Adam R., Guy L. F. Holburn, and Alessandro Piazza. “Activist Protest Spillovers into the Regulatory Domain: Theory and Evidence from the U.S. Nuclear Power Generation Industry.” *Organization Science* 33, no. 3 (May 2022): 1163–87. <https://doi.org/10.1287/orsc.2021.1473>.
- Friedman, Milton. “A Friedman Doctrine-- The Social Responsibility Of Business Is to Increase Its Profits.” *The New York Times*, September 13, 1970, sec. Archives. <https://www.nytimes.com/1970/09/13/archives/a-friedman-doctrine-the-social-responsibility-of-business-is-to.html>.
- Fuchs, Hailey. “Trump Attack on Diversity Training Has a Quick and Chilling Effect.” *The New York Times*, October 13, 2020, sec. U.S. <https://www.nytimes.com/2020/10/13/us/politics/trump-diversity-training-race.html>.
- Gelles, David. “Companies Stay Quiet on Texas’ New Abortion Law.” *The New York Times*, September 3, 2021, sec. Business. <https://www.nytimes.com/2021/09/03/business/companies-texas-sb8-abortion-law.html>.
- Hadani, Michael, and Douglas Schuler. “In Search of El Dorado: The Elusive Financial Returns on Corporate Political Investments.” *Strategic Management Journal* 34, no. 2 (2013): 165–81.
- Hambrick, Donald C., and Adam J. Wowak. “CEO Sociopolitical Activism: A Stakeholder Alignment Model.” *Academy of Management Review* 46, no. 1 (January 2021): 33–59. <https://doi.org/10.5465/amr.2018.0084>.

- Hebl, Michelle R., and Laura M. Mannix. "The Weight of Obesity in Evaluating Others: A Mere Proximity Effect." *Personality and Social Psychology Bulletin* 29, no. 1 (January 1, 2003): 28–38. <https://doi.org/10.1177/0146167202238369>.
- Hertel-Fernandez, Alexander. *Politics at Work: How Companies Turn Their Workers Into Lobbyists*. Oxford University Press, 2018.
- Holburn, Guy L. F., and Richard G. Vanden Bergh. "Integrated Market and Nonmarket Strategies: Political Campaign Contributions around Merger and Acquisition Events in the Energy Sector." *Strategic Management Journal* 35, no. 3 (2014): 450–60. <https://doi.org/10.1002/smj.2096>.
- Hou, Young, and Christopher W. Poliquin. "The Effects of CEO Activism: Partisan Consumer Behavior and Its Duration." *Strategic Management Journal* 44, no. 3 (2023): 672–703. <https://doi.org/10.1002/smj.3451>.
- Hurst, Reuben. "Strategically Ambiguous Identities: Evidence from Evangelical Clergy in Brazilian Elections." *The Journal of Politics* Forthcoming (2022).
- . "Workplace Backlash? Workforce Diversity and Dominant Group Attraction to Pro-Diversity Claims." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, February 20, 2021. <https://doi.org/10.2139/ssrn.3789682>.
- Hurst, Reuben, Saerom (Ronnie) Lee, and Justin Frake. "The Hidden Cost of Flat Hierarchies on Applicant Pool Diversity: Evidence from Experiments." SSRN Scholarly Paper. Rochester, NY, February 9, 2022. <https://doi.org/10.2139/ssrn.4030741>.
- Jan, Tracy, Jena McGregor, Renae Merle, and Natasha Tiku. "As Big Corporations Say 'Black Lives Matter,' Their Track Records Raise Skepticism." *Washington Post*, 2020. <https://www.washingtonpost.com/business/2020/06/13/after-years-marginalizing-black-employees-customers-corporate-america-says-black-lives-matter/>.
- Janney, Jay J., and Steve Gove. "Reputation and Corporate Social Responsibility Aberrations, Trends, and Hypocrisy: Reactions to Firm Choices in the Stock Option Backdating Scandal." *Journal of Management Studies* 48, no. 7 (2011): 1562–85. <https://doi.org/10.1111/j.1467-6486.2010.00984.x>.
- Jensen, Michael. "Should We Stay or Should We Go? Accountability, Status Anxiety, and Client Defections." *Administrative Science Quarterly* 51, no. 1 (2006): 97–128.
- Jonsson, Stefan, Henrich R. Greve, and Takako Fujiwara-Greve. "Undeserved Loss: The Spread of Legitimacy Loss to Innocent Organizations in Response to Reported Corporate Deviance." *Administrative Science Quarterly* 54, no. 2 (June 1, 2009): 195–228. <https://doi.org/10.2189/asqu.2009.54.2.195>.



- Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. “Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias.” *The Journal of Economic Perspectives* 5, no. 1 (1991): 193–206.
- Kalev, Alexandra, Frank Dobbin, and Erin Kelly. “Best Practices or Best Guesses? Assessing the Efficacy of Corporate Affirmative Action and Diversity Policies.” *American Sociological Review* 71, no. 4 (August 2006): 589–617. <https://doi.org/10.1177/000312240607100404>.
- Kang, Sonia K., Katherine A. DeCelles, András Tilcsik, and Sora Jun. “Whitened Résumés: Race and Self-Presentation in the Labor Market.” *Administrative Science Quarterly* 61, no. 3 (September 1, 2016): 469–502. <https://doi.org/10.1177/0001839216639577>.
- Khosrowshahi, Dara. “Being an Anti-Racist Company.” Uber Newsroom, July 17, 2020. <https://www.uber.com/newsroom/being-an-anti-racist-company/>.
- Krosnick, Jon A. “Government Policy and Citizen Passion: A Study of Issue Publics in Contemporary America.” *Political Behavior* 12, no. 1 (1990): 59–92.
- Larcker, David F, Stephen A Miles, and Brian Tayan. “The Double-Edged Sword of CEO Activism.” *Rock Center for Corporate Governance at Stanford University Closer Look Series: Topics, Issues and Controversies in Corporate Governance No. CGRP-74*, 2018, 12.
- Leonhardt, Megan. “Hard Right Turn: Companies Are Toeing the Political Line in Business-Friendly Red States.” *Fortune*, 2021. <https://fortune.com/2021/10/02/doing-business-in-conservative-states-texas-abortion-ban/>.
- Li, Zhao, and Richard W. Disalvo. “Can Stakeholders Mobilize Businesses for the Protection of Democracy? Evidence from the U.S. Capitol Insurrection.” *American Political Science Review*, October 3, 2022, 1–7. <https://doi.org/10.1017/S000305542200096X>.
- Luo, Jiao, Aseem Kaul, and Haram Seo. “Winning Us with Trifles: Adverse Selection in the Use of Philanthropy as Insurance.” *Strategic Management Journal* 39, no. 10 (October 2018): 2591–2617. <https://doi.org/10.1002/smj.2935>.
- MacFarquhar, Neil. “Four Years After ‘Unite the Right,’ Charlottesville Still Struggles to Move On.” *The New York Times*, November 18, 2021, sec. U.S. <https://www.nytimes.com/2021/11/18/us/charlottesville-nc-unite-the-right-rally.html>.
- Markus, Gregory B., and Philip E. Converse. “A Dynamic Simultaneous Equation Model of Electoral Choice.” *The American Political Science Review* 73, no. 4 (1979): 1055–70. <https://doi.org/10.2307/1953989>.
- Marquis, Christopher. “The Pressure of the Past: Network Imprinting in Intercorporate Communities.” *Administrative Science Quarterly* 48, no. 4 (2003): 655–89. <https://doi.org/10.2307/3556640>.

- McAuliffe, Terry. *Beyond Charlottesville: Taking a Stand Against White Nationalism*. St. Martin's Press, 2019.
- McConnell, Christopher, Yotam Margalit, Neil Malhotra, and Matthew Levendusky. "The Economic Consequences of Partisanship in a Polarized Era." *American Journal of Political Science* 62, no. 1 (January 1, 2018): 5–18. <https://doi.org/10.1111/ajps.12330>.
- McDonnell, Mary-Hunter, and Brayden King. "Keeping up Appearances: Reputational Threat and Impression Management after Social Movement Boycotts." *Administrative Science Quarterly* 58, no. 3 (September 1, 2013): 387–419. <https://doi.org/10.1177/0001839213500032>.
- McDonnell, Mary-Hunter, Kate Odziemkowska, and Elizabeth Pontikes. "Bad Company: Shifts in Social Activists' Tactics and Resources After Industry Crises." *Organization Science* 32, no. 4 (July 2021): 1033–55. <https://doi.org/10.1287/orsc.2020.1410>.
- McVeigh, Rory, and Kevin Estep. *The Politics of Losing: Trump, the Klan, and the Mainstreaming of Resentment*. Columbia University Press, 2019.
- Melloni, Gaia, Andrea Pataconi, and Nick Vikander. "Cashing in on the Culture Wars? CEO Activism, Wokewashing, and Firm Value." *Strategic Management Journal* n/a, no. n/a (2023). <https://doi.org/10.1002/smj.3542>.
- Mohliiver, Aharon, Donal Crilly, and Aseem Kaul. "Corporate Social Counterpositioning: How Attributes of Social Issues Influence Competitive Response." *Strategic Management Journal* 44, no. 5 (2023): 1199–1217. <https://doi.org/10.1002/smj.3461>.
- Naumovska, Ivana, and Dovev Lavie. "When an Industry Peer Is Accused of Financial Misconduct: Stigma versus Competition Effects on Non-Accused Firms." *Administrative Science Quarterly* 66, no. 4 (December 1, 2021): 1130–72. <https://doi.org/10.1177/00018392211020662>.
- Norris, Pippa, and Ronald Inglehart. *Cultural Backlash and the Rise of Populism: Trump, Brexit, and Authoritarian Populism*. Cambridge University Press, 2019.
- Panagopoulos, Costas, Donald P. Green, Jonathan Krasno, Michael Schwam-Baird, and Kyle Endres. "Partisan Consumerism: Experimental Tests of Consumer Reactions to Corporate Political Activity." *The Journal of Politics* 82, no. 3 (July 2020): 996–1007. <https://doi.org/10.1086/707491>.
- Parker, Kim, Nikki Graf, and Ruth Igielnik. "Generation Z Looks a Lot Like Millennials on Key Social and Political Issues." Pew Research Center's Social & Demographic Trends Project, January 17, 2019. <https://www.pewresearch.org/social-trends/2019/01/17/generation-z-looks-a-lot-like-millennials-on-key-social-and-political-issues/>.

- Penny, Helen, and Geoffrey Haddock. "Anti-Fat Prejudice among Children: The 'mere Proximity' Effect in 5-10 Year Olds." *Journal of Experimental Social Psychology* 43, no. 4 (2007): 678–83. <https://doi.org/10.1016/j.jesp.2006.07.002>.
- Piazza, Alessandro, and Julien Jourdan. "When the Dust Settles: The Consequences of Scandals for Organizational Competition." *Academy of Management Journal* 61, no. 1 (February 2018): 165–90. <https://doi.org/10.5465/amj.2015.1325>.
- Piazza, Alessandro, and Fabrizio Perretti. "Categorical Stigma and Firm Disengagement: Nuclear Power Generation in the United States, 1970–2000." *Organization Science* 26, no. 3 (June 2015): 724–42. <https://doi.org/10.1287/orsc.2014.0964>.
- Plaut, Victoria C., Flannery G. Garnett, Laura E. Buffardi, and Jeffrey Sanchez-Burks. "'What about Me?' Perceptions of Exclusion and Whites' Reactions to Multiculturalism." *Journal of Personality and Social Psychology* 101, no. 2 (2011): 337–53. <https://doi.org/10.1037/a0022832>.
- Pontikes, Elizabeth, Giacomo Negro, and Hayagreeva Rao. "Stained Red: A Study of Stigma by Association to Blacklisted Artists during the 'Red Scare' in Hollywood, 1945 to 1960." *American Sociological Review* 75, no. 3 (June 1, 2010): 456–78. <https://doi.org/10.1177/0003122410368929>.
- Pryor, John B., Glenn D. Reeder, and Andrew E. Monroe. "The Infection of Bad Company: Stigma by Association." *Journal of Personality and Social Psychology* 102, no. 2 (February 2012): 224–41. <https://doi.org/10.1037/a0026270>.
- Pullinger, Jenniefer. "Making Space: Using Diversity and Inclusion Programs to Build a Better Workplace." *C-Ville Weekly*, 2019. <https://www.c-ville.com/making-space-using-diversity-and-inclusion-programs-to-build-a-better-workplace/>.
- Reeve, Elle. "Charlottesville Civil Trial Will Explore Where Free Speech Becomes Conspiracy to Commit Violence." CNN, October 26, 2021. <https://www.cnn.com/2021/10/26/us/charlottesville-lawsuit-civil-trial/index.html>.
- "Representativeness Analysis of Lightcast Job Posting Data - U.S." Lightcast, 2022.
- Rosenhan, David, and Samuel Messick. "Affect and Expectation." *Journal of Personality and Social Psychology*, 1966, 38–44.
- Ross, Lee, David Greene, and Pamela House. "The 'False Consensus Effect': An Egocentric Bias in Social Perception and Attribution Processes." *Journal of Experimental Social Psychology* 13, no. 3 (May 1, 1977): 279–301. [https://doi.org/10.1016/0022-1031\(77\)90049-X](https://doi.org/10.1016/0022-1031(77)90049-X).

- Roth, Jonathan. “Pre-Test with Caution: Event-Study Estimates After Testing for Parallel Trends.” *Conditionally Accepted, American Economic Review: Insights*, 2021.
- Rozin, Paul, and Edward B. Royzman. “Negativity Bias, Negativity Dominance, and Contagion.” *Personality and Social Psychology Review* 5, no. 4 (November 1, 2001): 296–320. [https://doi.org/10.1207/S15327957PSPR0504\\_2](https://doi.org/10.1207/S15327957PSPR0504_2).
- Short, Jodi L., and Michael W. Toffel. “Making Self-Regulation More Than Merely Symbolic: The Critical Role of the Legal Environment.” *Administrative Science Quarterly* 55, no. 3 (September 1, 2010): 361–96. <https://doi.org/10.2189/asqu.2010.55.3.361>.
- Signer, Michael. *Cry Havoc: Charlottesville and American Democracy Under Siege*. Public Affairs, 2020.
- The Economist*. “In the Trump Era, Big Business Is Becoming More Political.” November 30, 2017. <https://www.economist.com/business/2017/11/30/in-the-trump-era-big-business-is-becoming-more-political>.
- Thompson, Carolyn, and Heahter Hollingsworth Press. “Some School Systems Pause Diversity Programs amid Pushback.” *Los Angeles Times*, February 20, 2022, sec. World & Nation. <https://www.latimes.com/world-nation/story/2022-02-20/some-school-systems-pause-diversity-programs-amid-pushback>.
- Tolbert, Pamela S., and Emilio J. Castilla. “Editorial Essay: Introduction to a Special Issue on Inequality in the Workplace (‘What Works?’).” *ILR Review* 70, no. 1 (January 1, 2017): 3–15. <https://doi.org/10.1177/0019793916675336>.
- Tomz, Michael, and Robert Van Houweling. “The Electoral Implications of Candidate Ambiguity.” *American Political Science Review* 103, no. 1 (February 2009): 83–98. <https://doi.org/10.1017/S0003055409090066>.
- UVA Center for Politics. “New Poll: Some Americans Express Troubling Racial Attitudes Even as Majority Oppose White Supremacists,” September 14, 2017. <https://centerforpolitics.org/crystalball/articles/new-poll-some-americans-express-troubling-racial-attitudes-even-as-majority-oppose-white-supremacists/>.
- Vogel, David. *The Market for Virtue: The Potential and Limits of Corporate Social Responsibility*. Brookings Institution Press, 2005.
- Wagner, Tillmann, Richard J. Lutz, and Barton A. Weitz. “Corporate Hypocrisy: Overcoming the Threat of Inconsistent Corporate Social Responsibility Perceptions.” *Journal of Marketing* 73, no. 6 (November 1, 2009): 77–91. <https://doi.org/10.1509/jmkg.73.6.77>.
- Werner, Timothy. “Investor Reaction to Covert Corporate Political Activity.” *Strategic Management Journal* 38, no. 12 (2017): 2424–43. <https://doi.org/10.1002/smj.2682>.

———. *Public Forces and Private Politics in American Big Business*. Cambridge University Press, 2012.

“Wide Gender Gap, Growing Educational Divide in Voters’ Party Identification.” Pew Research Center, 2018. <https://www.pewresearch.org/politics/2018/03/20/1-trends-in-party-affiliation-among-demographic-groups/>.

Wowak, Adam J., John R. Busenbark, and Donald C. Hambrick. “How Do Employees React When Their CEO Speaks Out? Intra- and Extra-Firm Implications of CEO Sociopolitical Activism.” *Administrative Science Quarterly*, February 17, 2022. <https://doi.org/10.1177/00018392221078584>.

Yue, Lori Qingyuan, Hayagreeva Rao, and Paul Ingram. “Information Spillovers from Protests against Corporations: A Tale of Walmart and Target.” *Administrative Science Quarterly* 58, no. 4 (December 1, 2013): 669–701. <https://doi.org/10.1177/0001839213511243>.

## 2.9 Tables

Table 2.9.1 Summary Statistics

	(1) Full Sample	(2) Charlottesville MSA	(3) Flagship Counties	(4) Clinton Counties
<i>Job-level</i>				
Pro-diversity claim	0.12 (0.32)	0.15 (0.36)	0.13 (0.34)	0.13 (0.33)
Equal opportunity claim	0.40 (0.49)	0.49 (0.50)	0.41 (0.49)	0.40 (0.49)
Job requires bachelor's degree	0.29 (0.45)	0.39 (0.49)	0.32 (0.47)	0.34 (0.47)
Pay annually dummy	0.15 (0.35)	0.16 (0.37)	0.17 (0.38)	0.15 (0.36)
Moving expenses covered	0.05 (0.22)	0.03 (0.17)	0.05 (0.22)	0.05 (0.21)
Postings	66,498,855	62,074	6,080,342	42,044,363
<i>Employer-level</i>				
Annual jobs / employer	29.05 (686.84)	11.78 (104.07)	17.02 (178.94)	25.10 (502.24)
Common employers (share)	Anthem (0.01) Lowe's (0.01) CACI (0.01) Marriott (0.01) HCA (<0.01)	UVA (0.19) CACI (0.10) Anthem (0.05) Sentara (0.02) UVA Health (0.02)	Amazon (0.01) Anthem (0.01) JP Morgan Chase (0.01) Lowe's (0.01) Accenture (0.01)	Anthem (0.01) CACI (0.01) Marriott (0.01) Accenture (0.01) Macy's (0.01)
Common industries (share)	Healthcare (0.19) Finance (0.10) Professional Ser. (0.10) Hospitality (0.09) Wholesale (0.09)	Education (0.23) Professional Ser. (0.15) Healthcare (0.13) Finance (0.09) Retail (0.08)	Healthcare (0.17) Professional Ser. (0.10) Education (0.10) Finance (0.10) Hospitality (0.08)	Healthcare (0.18) Professional Ser. (0.12) Finance (0.11) Hospitality (0.10) Retail (0.08)
Public administration	0.03 (0.17)	0.04 (0.19)	0.03 (0.18)	0.03 (0.16)
Unique employers	1,493,729	3,167	230,395	1,096,955
Annual jobs / employer-county	8.45 (62.04)	10.32 (97.13)	9.48 (81.51)	9.45 (74.77)
Unique employer-county dyads	5,034,136	3,708	407,585	2,874,805
<i>County-level</i>				
Median age	41.19 (5.33)	40.77 (5.77)	34.90 (4.77)	37.93 (5.13)
% White	77.03 (19.84)	74.64 (8.42)	70.23 (18.27)	50.39 (23.80)
Population	102,976.49 (32,9057.10)	38,137.37 (3,2325.66)	398,565.19 (444,815.83)	392,209.98 (777,446.08)
Average annual earnings	33,698.15 (6,334.39)	37,580.10 (5,695.84)	39,330.05 (6,117.48)	35,955.93 (9,847.74)
Clinton vote share in 2016	0.32 (0.15)	0.50 (0.16)	0.55 (0.13)	0.61 (0.09)
Unique counties	3,141	6	54	432

Table 2.9.2 Difference-in-Differences Estimates at the Level of Job Postings

Sample	(1) All	(2) Flagship Counties	(3) Clinton Counties	(4) Only UVA	(5) Excluding UVA
Charlottesville × Post-rally	0.086*** (0.003)	0.086*** (0.007)	0.089*** (0.004)	0.215*** (0.036)	0.024*** (0.005)
Year-month fixed effects	✓	✓	✓	✓	✓
Employer-county fixed effects	✓	✓	✓	✓	✓
Job-level controls	✓	✓	✓	✓	✓
County-level controls	✓	✓	✓	✓	✓
Observations	55,668,345	5,089,867	34,632,156	1,273,264	55,654,371

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

\* All estimates are linear probability models with robust standard errors clustered on county. Notice that the discrepancy in sample size in this table compared to Column 1 of Table 1 results from the fact that I estimated these models with Stata's *reghdfe* command, which drops singleton observations.

Table 2.9.3 Alternative Explanations and Accompanying Empirical Analysis

	Alternative Explanation	Patterns Inconsistent with Explanation
Isomorphism	The effect chiefly reflected a general imitative process or an emergent norm of making pro-diversity claims that diffused across Charlottesville's employers.	Adoption of pro-diversity claims was not exceptionally widespread but was concentrated in instances when firms faced greater incentives to do so, including when job seekers were more likely to presume that employers held anti-diversity positions and when job seekers held strong pro-diversity preferences.
Issue salience	The effect chiefly reflected the fact that issues of diversity were suddenly very salient to Charlottesville's workers, incentivizing employers to appeal to these preferences.	The effect was larger when targeting job seekers outside of Charlottesville, and the effect emerged after a delay and then grew rather than appearing suddenly, which would be more consistent with a sudden spike in salience.
Upper echelon	The effect chiefly reflected the personal values of top managers and/or the fact that top managers were suddenly paying greater attention to diversity issues.	Holding constant top managers, the increase in diversity claims was greater in Charlottesville compared to other firm locations. The delayed effect, supported by supplementary interviews, is consistent with the idea that increased use of pro-diversity claims reflected internal deliberation regarding strategic implications of stance taking rather than unilateral expression of top managers' values.
EO compliance	The effect was chiefly a byproduct of firms' efforts to signal compliance with the EOA by including EO claims in their job postings. Because EO claims often include pro-diversity claims, there was a mechanical increase in pro-diversity claims.	The increase in EO claims was much smaller in magnitude than the increase in pro-diversity claims.

Table 2.9.4 Adoption of Pro-Diversity Claims Following the Rally

	(1)	(2)	(3)
Charlottesville × Post-rally × Charlottesville headquarters	0.176*** (0.005)		
Charlottesville × Post-rally × Requires bachelor's degree		0.028*** (0.002)	
Charlottesville × Post-rally × Moving expenses covered			0.165*** (0.009)
Year-month fixed effects	✓	✓	✓
Employer-county fixed effects	✓	✓	✓
Job-level controls	✓	✓	✓
County-level controls	✓	✓	✓
Observations	55,668,345	55,668,345	55,668,345

\*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ .  
 \*For brevity, the table excludes the main effects and shows only the estimated interaction effect ( $\hat{\beta}_1$ ).

Table 2.9.5 Additional Analyses Testing Theorized Mechanism

<b>Outcome</b>	(1) <b>Pro-Diversity</b>	(2) <b>EO Claim</b>	(3) <b>“Charlottesville”</b>
Charlottesville × Post-rally	0.010** (0.004)	0.035*** (0.004)	-0.018*** (0.001)
Year-month fixed effects	✓	✓	✓
Employer-county fixed effects	✓	✓	✓
Employer-year-month fixed effects job-level controls	✓	✓	✓
County-level controls	✓	✓	✓
Observations	53,921,070	55,668,345	55,668,345

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



Table 2.9.6 Relationship Between Unite the Right Rally and Newfound Wage Premiums in Job Postings

<b>Outcome</b>	<b>(1) Annual Salary</b>	<b>(2) Log Annual Salary</b>	<b>(3) Annual Salary</b>	<b>(4) Log Annual Salary</b>
Charlottesville × Post-rally	2500.027*** (503.652)	0.038*** (0.007)	2564.113*** (527.743)	.045*** (0.007)
Pro-diversity claim			3126.015*** (338.612)	0.054*** (0.004)
Charlottesville × Pro-diversity claim			-2827.620*** (499.297)	-0.063*** (0.006)
Year-month fixed effects	✓	✓	✓	✓
Employer-county fixed effects	✓	✓	✓	✓
Job-level controls	✓	✓	✓	✓
County-level controls	✓	✓	✓	✓
Observations	7,028,116	7,028,116	7,028,116	7,028,116

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

\* The outcome in Models 1 and 3 is the minimum annual salary advertised in the job posting. The outcome in Models 2 and 4 is the log of the minimum annual salary advertised in the job posting. All models are estimated using ordinary least squares with robust standard errors clustered on county. The relatively low sample size in these models follows from the fact that many jobs do not list annual salaries.

## 2.10 Figures

Figure 2.10.1 Association Between Charlottesville and White Supremacy After the Unite the Right Rally

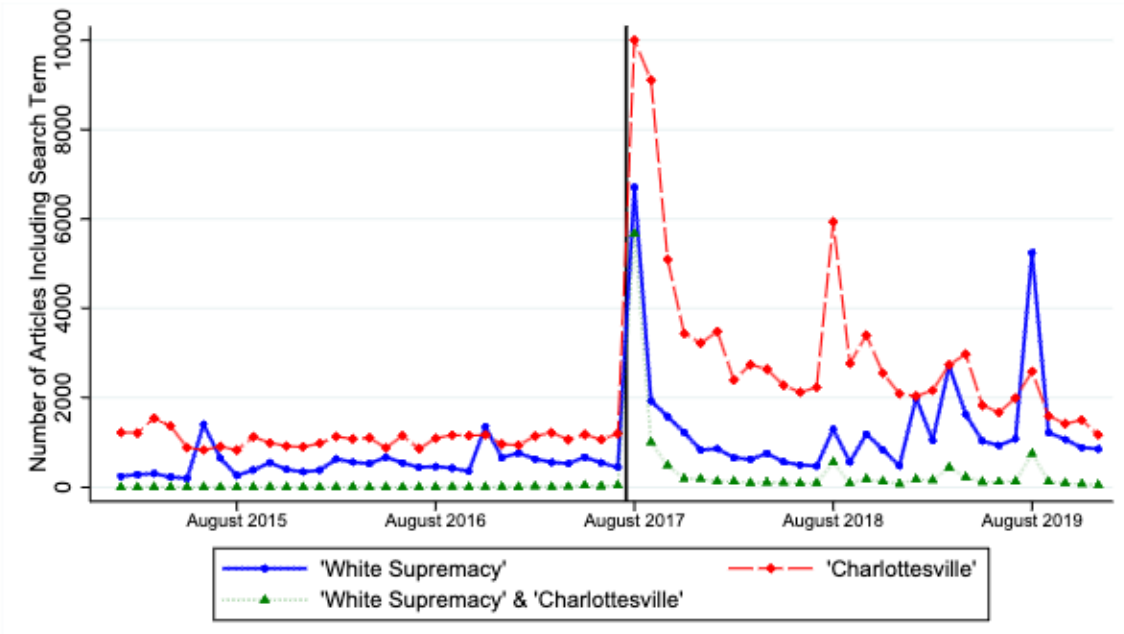
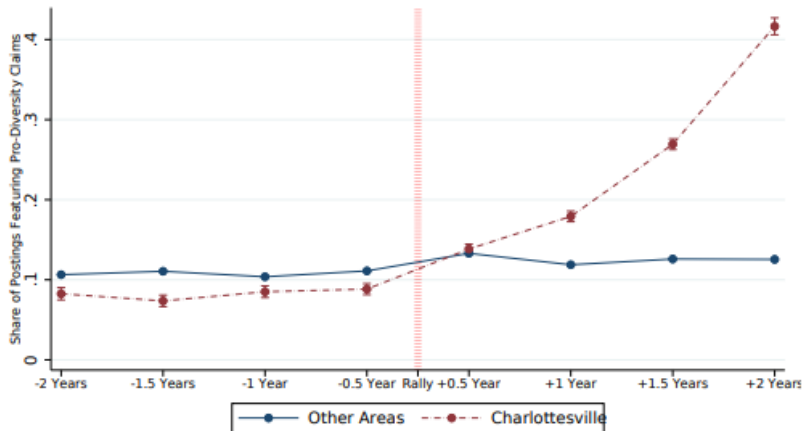
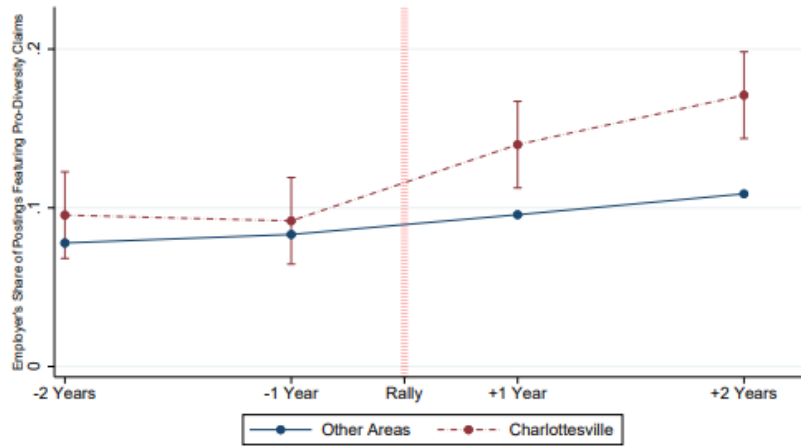


Figure 2.10.2 Share of Charlottesville Job Postings Featuring Pro-Diversity Claims



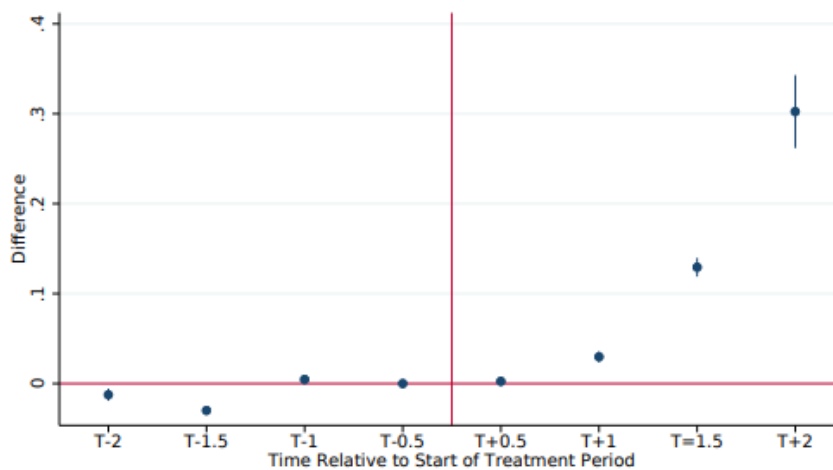
NOTES: Points indicate the share of job postings by six month periods that included pro-diversity claims in the Charlottesville MSA versus other areas, (i.e., the rest of the United States). Six months periods are measured relative to the second day of the rally, August 12, 2017. The vertical dashed line marks the division between the pre and post-rally period. Points are accompanied by 95% confidence intervals, but given high precision these are not visible on the "Other Areas" points are only barely visible on the Charlottesville points.

Figure 2.10.3 Share of Job Postings in Which Charlottesville Employers Featured Pro-Diversity Claims



NOTES: Points indicate the share of job postings, on average, in which employer-county dyads featured pro-diversity claims in their job postings in the Charlottesville MSA versus other areas, (i.e., the rest of the United States). This reflects a balanced panel in the sense that it only includes employer-county dyads that post at least one job in each of the four years surrounding the rally. Years are measured relative to the second day of the rally, August 12, 2017. The vertical dashed line marks the division between the pre and post-rally period. Bars represent 95% confidence intervals, but given high precision these are not visible on the “Other Areas” points.

Figure 2.10.4 Dynamic Difference-in-Differences Estimates of  $\delta_k$  from Equation 2



NOTES: Each point represents the estimates of  $\delta_k$  from Equation 2, where  $\delta_{-3}$  corresponds to “T-2”,  $\delta_{-2}$  corresponds to “T-1.5,” and so on. These represent the difference in predicted probability that, in the indicated time period, a job posting in Charlottesville versus a job posting outside of Charlottesville features a pro-diversity claim relative to this difference in the reference time period ( $\delta_0$ , which corresponds to “T-0.5”). The reference time period is the six-month period preceding the Unite the Right Rally (February 12, 2017 - August 11, 2017). Bars indicate 95% confidence intervals.

## 2.11 Appendices

Table 2.11.1 Estimated effects using employer, rather than employer-county, fixed effects.

Sample:	(1) All	(2) Flagship Counties	(3) Clinton Counties	(4) Only UVA	(5) Excluding UVA
Charlottesville × Post-Rally	0.036*** (0.007)	0.044*** (0.007)	0.039*** (0.004)	0.169*** (0.047)	0.012*** (0.002)
Year Month Fixed Effects	✓	✓	✓	✓	✓
Employer Fixed Effects	✓	✓	✓	✓	✓
Job-Level Controls	✓	✓	✓	✓	✓
County-Level Controls	✓	✓	✓	✓	✓
Observations	56,900,934	5,149,846	35,254,610	1,279,906	56,886,915

Notes: The units of analysis in all models are job postings. “Charlottesville × Post-Rally” is a dummy variable equal to one for job postings made in the Charlottesville area after August 11, 2017. The outcome in all models is a dummy variable equal to one if the job posting included a pro-diversity claim. All columns report estimates of a version of Equation 1 that includes employer level rather than employer-county level fixed effect and differ only in terms of the sample on which they are estimated. Model 1 includes the full sample of job postings. Model 2 includes the subset of job postings in the Charlottesville area and counties that were also home to flagship state universities. Model 3 includes the subset of job postings in counties where more than 50% of the 2016 presidential vote went to Hilary Clinton. Model 4 includes job postings from the University of Virginia and all other job postings by institutions of higher education outside of Charlottesville. Model 5 excludes job postings made by the University of Virginia. The discrepancy in sample size in this table compared to Model 1 of Table 1 follows from the fact that I estimate these model’s with Stata’s *reghdfe* command, which drops singleton observations. † $p < 0.10$ , \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$

Table 2.11.2 Summary Statistics for Balanced-by-Year Panel of Employee-County Dyads

	(1) Full Sample	(2) Charlottesville MSA	(3) Flagship Counties	(4) Clinton Counties
Pro-Diversity Claims	0.09 (0.25)	0.12 (0.29)	0.09 (0.25)	0.09 (0.24)
Equal Opportunity Claims	0.40 (0.43)	0.49 (0.45)	0.40 (0.43)	0.39 (0.43)
Number of Jobs	28.60 (134.81)	31.96 (199.53)	33.26 (177.89)	34.51 (165.41)
Share of Jobs Requiring BA	0.23 (0.27)	0.24 (0.29)	0.25 (0.27)	0.29 (0.29)
Share of Jobs with Moving Expenses Covered	0.06 (0.16)	0.05 (0.16)	0.06 (0.16)	0.06 (0.15)
Unique Employers	72,734	304	13,447	55,692
Unique Employer-County Dyads	390,942	322	32,142	215,023
Observations (Employer-County-Year)	1,563,768	1,288	128,568	860,092

Notes: The units of analysis are employer-county-years. These units comprise a panel of employer-county dyads balanced on the four years surrounding the treatment event.

Table 2.11.3 Difference-in-differences estimates at the employer-county level show a significant increase in pro-diversity claims among Charlottesville job postings following the Unite the Right rally.

Sample:	(1) All	(2) Flagship Counties	(3) Clinton Counties	(4) Only UVA	(5) Excluding UVA
Charlottesville × Post-Rally	0.039* (0.018)	0.043* (0.018)	0.028*** (0.002)	0.824*** (0.029)	0.039* (0.018)
Year Fixed Effects	✓	✓	✓	✓	✓
Employer-County Fixed Effects	✓	✓	✓	✓	✓
Observations	1,553,124	128,568	849,448	7,020	1,553,084

*Notes:* The units of analysis in all models are employer-county-years. These units comprise a panel of employer-county dyads balanced on the four years surrounding the treatment event. Years are measured relative to the treatment event. The outcome in all models is the share of job postings by a given employer in a given county in a given year that feature a pro-diversity claim. “Charlottesville × Post-Rally” is a dummy variable equal to one for employers in any of the counties comprising the Charlottesville MSA after August 11, 2017. All columns report estimates of Equation 1 and differ only in terms of the sample on which they are estimated. Model 1 includes the full sample of job postings. Model 2 includes the subset of job postings in the Charlottesville MSA and counties that were also home to flagship state universities. Model 3 includes the subset of job postings in the Charlottesville MSA and counties where more than 50% of the 2016 presidential vote went to Hilary Clinton. Model 4 includes job postings from the University of Virginia and all other job postings by institutions of higher education outside of Charlottesville. Model 5 excludes job postings made by the University of Virginia. All estimates are linear probability models with robust standard errors clustered on county. † $p < 0.10$ , \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

Table 2.11.4 Analyzing at the employer level, the estimated effect grew in the post-treatment period.

	(1)
Charlottesville × 2 Years Prior	0.009*** (0.001)
Charlottesville × 1 Year Following	0.035 (0.020)
Charlottesville × 2 Years Following	0.053** (0.016)
Year Fixed Effects	✓
Employer-County Fixed Effects	✓
Observations	1,553,124

*Notes:* The units of analysis are employer-county-years. These units comprise a panel of employer-county dyads balanced on the four years surrounding the treatment event. Years are measured relative to the treatment event. The outcome is the share of job postings by a given employer in a given county in a given year that feature a pro-diversity claim. “Charlottesville × Post-Rally” is a dummy variable equal to one for employers in any of the counties comprising the Charlottesville MSA after August 11, 2017. This table reports estimates of a version of Equation 1 where *Charlottesville* is instead interacted with dummies for each of the four years comprising the study period. The reference year is the year leading up to the rally. I estimate the equation with a linear probability model with robust standard errors clustered on county. † $p < 0.10$ , \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

Table 2.11.5 Consistent with the proposed mechanism, employer-level analysis shows that adoption of pro-diversity claims was larger where targeted job seekers were more likely to generalize the rally and held stronger pro-diversity preferences.

	(1)	(2)	(3)
Charlottesville × Post-Rally × Charlottesville Headquarters	0.056** (0.020)		
Charlottesville × Post-Rally × Employer Dependent on Bachelor's-Holding Workers		0.046*** (0.014)	
Charlottesville × Post-Rally × Employer Reliant on Out of Town Workers			0.242 <sup>†</sup> (0.143)
Year Fixed Effects	✓	✓	✓
Employer-County Fixed Effects	✓	✓	✓
Observations	1,553,124	1,553,124	1,553,124

*Notes:* The units of analysis in all models are employer-county-years. These units comprise a panel of employer-county dyads balanced on the four years surrounding the treatment event. Years are measured relative to the treatment event. The outcome in all models is the share of job postings by a given employer in a given county in a given year that feature a pro-diversity claim. “Charlottesville × Post-Rally” is a dummy variable equal to one for employers in any of the counties comprising the Charlottesville MSA after August 11, 2017. All columns report estimates of a version of Equation 3 that includes an interaction term and differ only in terms of which variable is interacted. For brevity, I exclude from the table the main effects and only show the estimated interactions of interest. “Charlottesville Headquarters” is a dummy variable equal to one for job postings made in the Charlottesville area after August 11, 2017. “Employer Dependent on Bachelor's-Holding Workers” is a dummy equal to one for employers that are in the top 75% in terms of the total number of postings in the study period that required at least a Bachelor's degree. “Employer Reliant on Out of Town Workers” is a dummy equal to one for employers that, across the study period, offered to cover moving expenses in over half of their job postings. <sup>†</sup> $p < 0.10$ , \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

Table 2.11.6 Additional employer-level analysis is consistent with the proposed mechanism.

Outcome:	(1) Pro-Diversity	(2) Equal Opportunity	(3) “Charlottesville”
Charlottesville × Post-Rally	0.011 (0.008)	0.020 (0.012)	-0.042*** (0.002)
Year Month Fixed Effects	✓	✓	✓
Employer County Fixed Effects	✓	✓	✓
Employer-Year-Month Fixed Effects	✓	✓	✓
Observations	1,343,940	1,553,124	1,553,124

*Notes:* The units of analysis in all models are employer-county-years. These units comprise a panel of employer-county dyads balanced on the four years surrounding the treatment event. Years are measured relative to the treatment event. “Charlottesville × Post-Rally” is a dummy variable equal to one for employers in Charlottesville after August 11, 2017. Model 1 is identical to Equation 3 with the difference that it includes employer-year-month fixed effects. The outcome in Model 1 is the share of job postings by a given employer in a given county in a given year that feature a pro-diversity claim. The outcome in Model 2 is the share of job postings by a given employer in a given county in a given year that feature an equal opportunity claim. The outcome in Model 3 is the share of job postings by a given employer in a given county in a given year that feature the word “Charlottesville.” Models 2 and 3 differ from Equation 3 only in terms of the outcome variable. <sup>†</sup> $p < 0.10$ , \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

Table 2.11.7 The synthetic control method achieves better balance on observables compared to the full sample of MSAs.

	Charlottesville MSA	Synthetic Charlottesville MSA	All MSAs
Median Age	31	31	37
Population	44,751	75,234	377,849
Median Income (\$)	36,886	36,880	36,583
% White	66	66	69
Number of Jobs Per Month	1,250	1,209	3,247

Notes: The units of analysis are MSA-months where months are measured relative to the treatment date.

Table 2.11.8 Excluding postings by UVA, the rally was associated with a newfound wage premium in Charlottesville job postings, but this premium was lower among job postings that featured pro-diversity claims.

Outcome:	(1) Annual Salary	(2) Log Annual Salary	(3) Annual Salary	(4) Log Annual Salary
Charlottesville × Post-Rally	1867.845* (835.387)	0.040*** (0.012)	2116.398* (880.903)	0.044*** (0.012)
Pro-Diversity Claim			3129.730*** (338.735)	0.054*** (0.004)
Charlottesville × Pro-Diversity Claim			-2978.107*** (434.749)	-0.057*** (0.007)
Year-Month Fixed Effects	✓	✓	✓	✓
Employer-County Fixed Effects	✓	✓	✓	✓
Job-Level Controls	✓	✓	✓	✓
County-Level Controls	✓	✓	✓	✓
Observations	7,022,598	7,022,598	7,022,598	7,022,598

Notes: This table differs from Table 6 only in that it excludes postings by UVA. The units of analysis in all models are job posting. "Charlottesville × Post-Rally" is a dummy variable equal to one for job postings made in the Charlottesville area after August 11, 2017. The outcome in models 1 and 3 is the minimum annual salary advertised in the job posting. The outcome in Models 2 and 4 is the log of the minimum annual salary advertised in the job posting. All models are estimated using ordinary least squares with robust standard errors clustered on county. Notice that the relatively low sample size in these models follows from the fact that many jobs do not list annual salaries. <sup>†</sup> $p < 0.10$ , \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

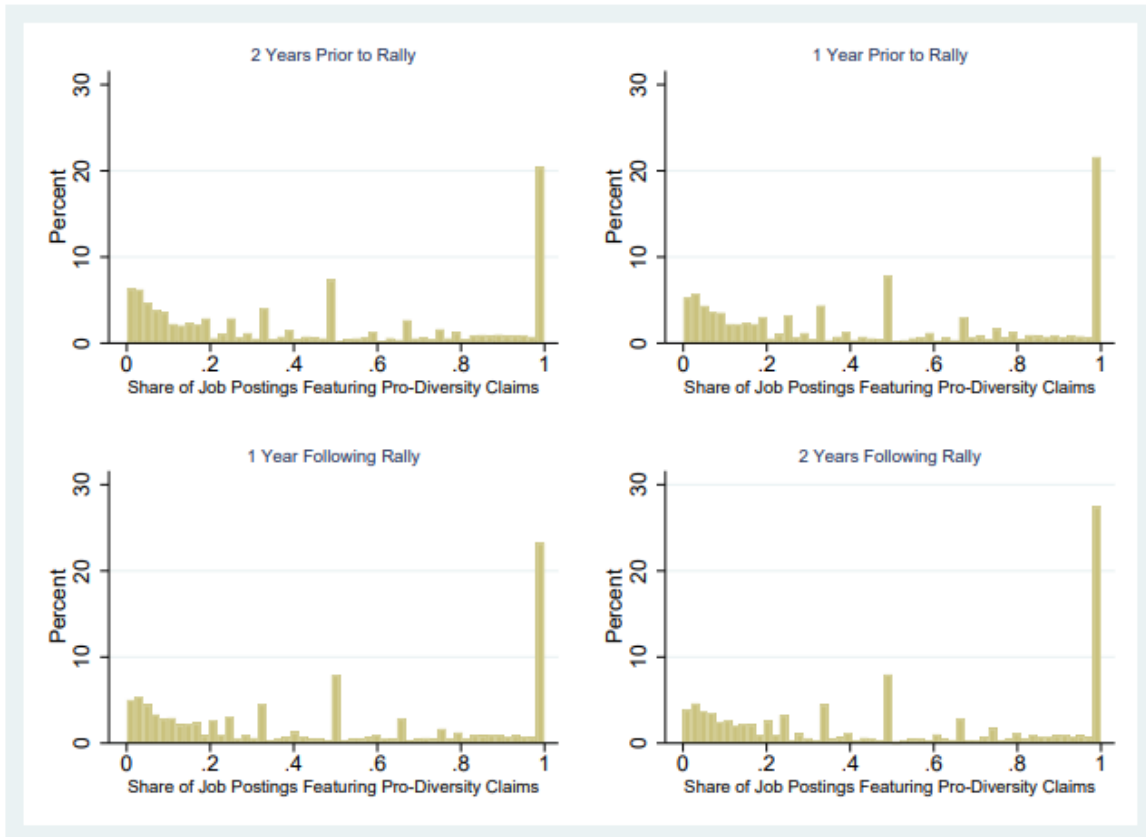
Figure 2.11.1 A photo from the rally illustrates that pro-diversity claims represented a sociopolitical stance in opposition to the anti-diversity positions of white supremacists.



*NOTES:* This photo, featured in Reeve (2021), was taken in Charlottesville during the the Unite the Right Rally. It shows a prominently displayed banner bearing the pro-diversity claim “Diversity makes us stronger” above agitators bearing flags with anti-diversity symbols.

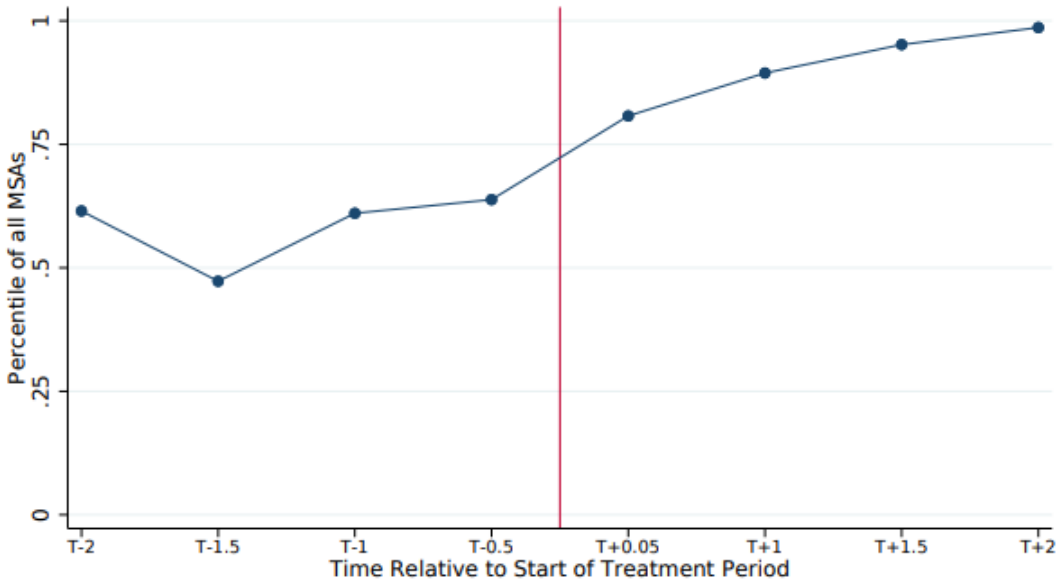


Figure 2.11.2 Among employer-county dyads that feature a pro-diversity claims in any of their job postings in a given year, the percentage of employer-county dyads that feature pro-diversity claims in all of their job postings never rose above 30%.




NOTES: This chart shows that, in most cases, employers that feature pro-diversity claims in their job postings do not do so in *all* of their job postings. Each of the four charts are histograms showing the percentage of employer-county dyads in the indicated year that featured pro-diversity claims in a given share of their job postings. Only employers that post more than one job in a given year are included, since only these can possibly vary their use of pro-diversity claims across postings. The sample is limited to employers that featured a pro-diversity claim in at least one job posting in that year. The purpose of this figure is to show that there is considerable variation in the extent to which employers feature pro-diversity claims in their job postings. If employers mostly make one-off decisions to include pro-diversity claims in *all* of their job postings, then all or most of the mass in these figures would be at 1 (*i.e.*, the far-right of the graph). Instead, the mass at 1 never rises above 30% and in the first three years is below 25%.

Figure 2.11.3 Following the rally, the Charlottesville MSA rose to the 99th percentile of all MSAs in terms of job postings featuring pro-diversity claims.




NOTES: Each point indicates, for the given 6-month period, the Charlottesville MSA's percentile relative to all other MSAs in terms of the share of job postings that featured pro-diversity claims.

Figure 2.11.4 The share of job postings made by the University of Virginia that featured pro-diversity claims increased significantly following the rally.




NOTES: Points indicate the fraction of job postings by year that included pro-diversity claims for the University of Virginia versus other higher education institutions in the United States. The vertical dashed line marks the division between the pre and post-rally period. Points are accompanied by 95% confidence intervals, but given high precision these are not visible on the “Other Higher Education Institutions” points.

Figure 2.11.5 Treatment effect at the true treatment threshold is significantly larger than treatment effects estimates at placebo thresholds six months, one year, and eighteen months before treatment.




*NOTES:* Points are estimates of  $\beta$  from Equation 1, which differ in terms of the definition of the treatment threshold and sample of job postings. The first three points are for the placebo treatment thresholds labelled on the x-axis and exclude all observations in the post-shock period (in other words, include only those job postings in the two years leading up to the rally). The estimate on the far right, marked in red, is the estimated treatment effect with the full sample at the true treatment threshold.

Figure 2.11.6 Excluding postings by UVA, dynamic difference-in-difference estimates show the estimated effect grew throughout the post-treatment period.



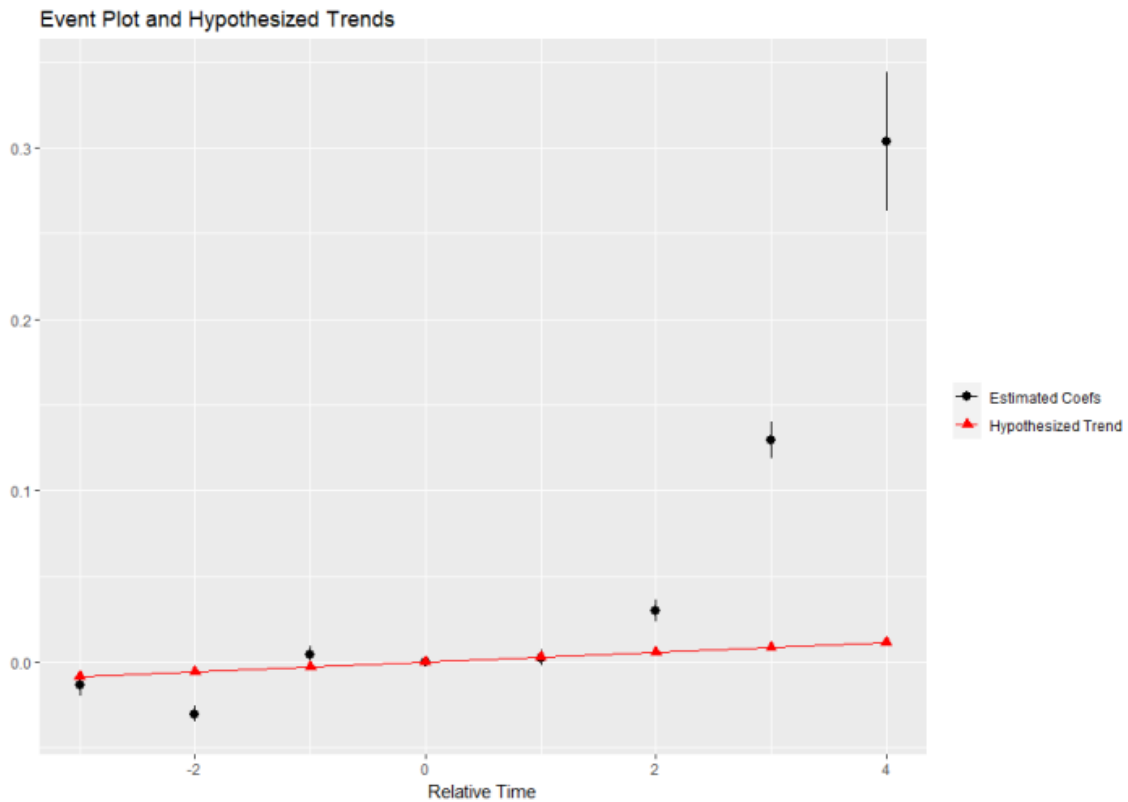
*NOTES:* This figure is analogous to Figure 4 but excludes postings by UVA. Each point represents the estimates of  $\delta_k$  from Equation 2, which represent the difference in predicted probability that, in the indicated time period, a job posting in Charlottesville versus a job posting outside of Charlottesville features a pro-diversity claim relative to this difference in the reference time period ( $\delta_0$ ). The reference time period is the six month period preceding the Unite the Right Rally (February 12, 2017 - August 11, 2017). Bars indicate 95% confidence intervals. The region to the right of the vertical center line indicates the post-shock period.

Figure 2.11.7 Robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 50% power level.




*NOTES:* The circular points and associated confidence intervals are the same time-specific difference-in-differences estimates presented in Figure 4. The triangular points and accompanying trend line indicates the hypothetical non-parallel pre-trend that would be detected at a power level of 50%. This is generated using the methods and software presented in Roth (2021). This line does not cross the 95% confidence intervals of the positive effects estimated in the post-treatment period ( $\delta_2$ ,  $\delta_3$ , and  $\delta_4$ , or the final three points on the right), which is consistent with the conclusion that these estimates are robust to a violation of parallel pre-trends of this size.

Figure 2.11.8 Robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 80% power level.



*NOTES:* The circular points and associated confidence intervals are the same time-specific difference-in-differences estimates presented in Figure 4. The triangular points and accompanying trend line indicates the hypothetical non-parallel pre-trend that would be detected at a power level of 80%. This is generated using the methods and software presented in Roth (2021). This line does not cross the 95% confidence intervals of the positive effects estimated in the post-treatment period ( $\delta_2$ ,  $\delta_3$ , and  $\delta_4$ , or the final three points on the right), which is consistent with the conclusion that these estimates are robust to a violation of parallel pre-trends of this size.

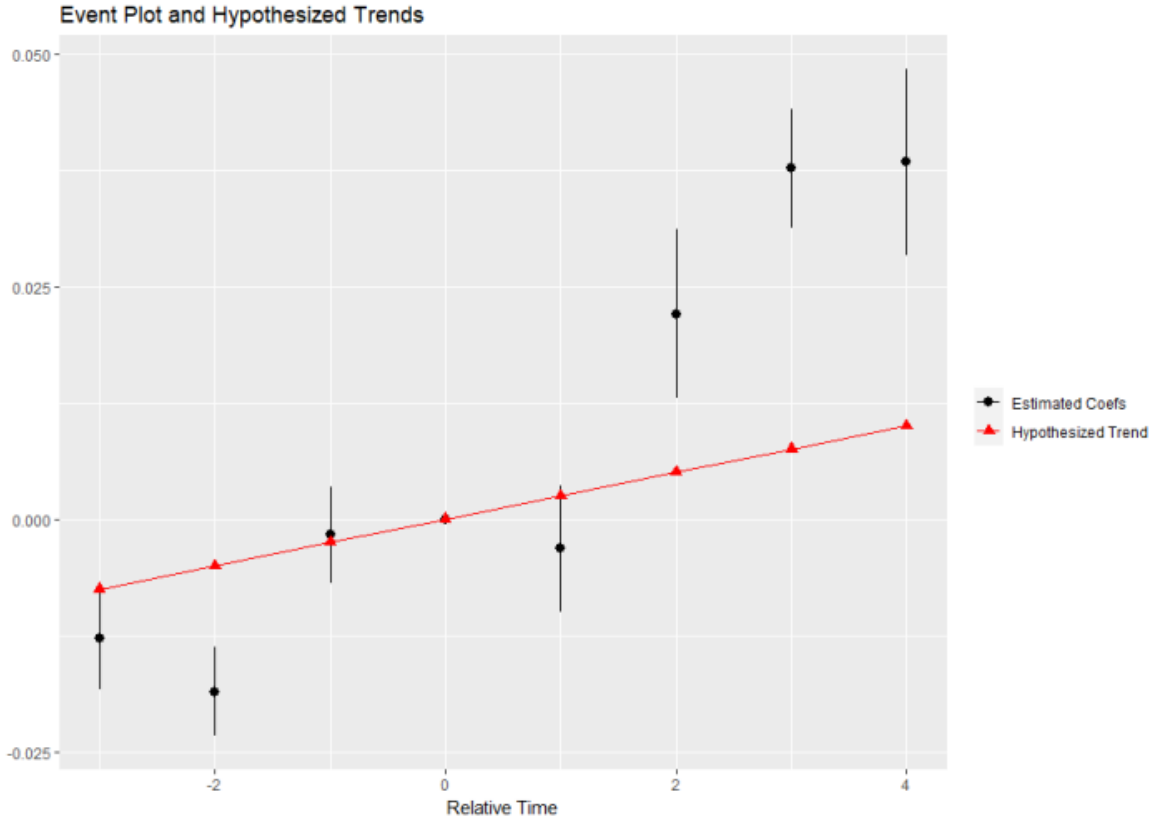
Figure 2.11.9 Excluding postings by UVA, robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 50% power level.



NOTES: The circular points and associated confidence intervals are the same time-specific difference-in-differences estimates presented in Figure A6. The triangular points and accompanying trend line indicates the hypothetical non-parallel pre-trend that would be detected at a power level of 50%. This is generated using the methods and software presented in Roth (2021). This line does not cross the 95% confidence intervals of the positive effects estimated in the post-treatment period ( $\hat{\delta}_2$ ,  $\hat{\delta}_3$ , and  $\hat{\delta}_4$ , or the final three points on the right), which is consistent with the conclusion that these estimates are robust to a violation of parallel pre-trends of this size.

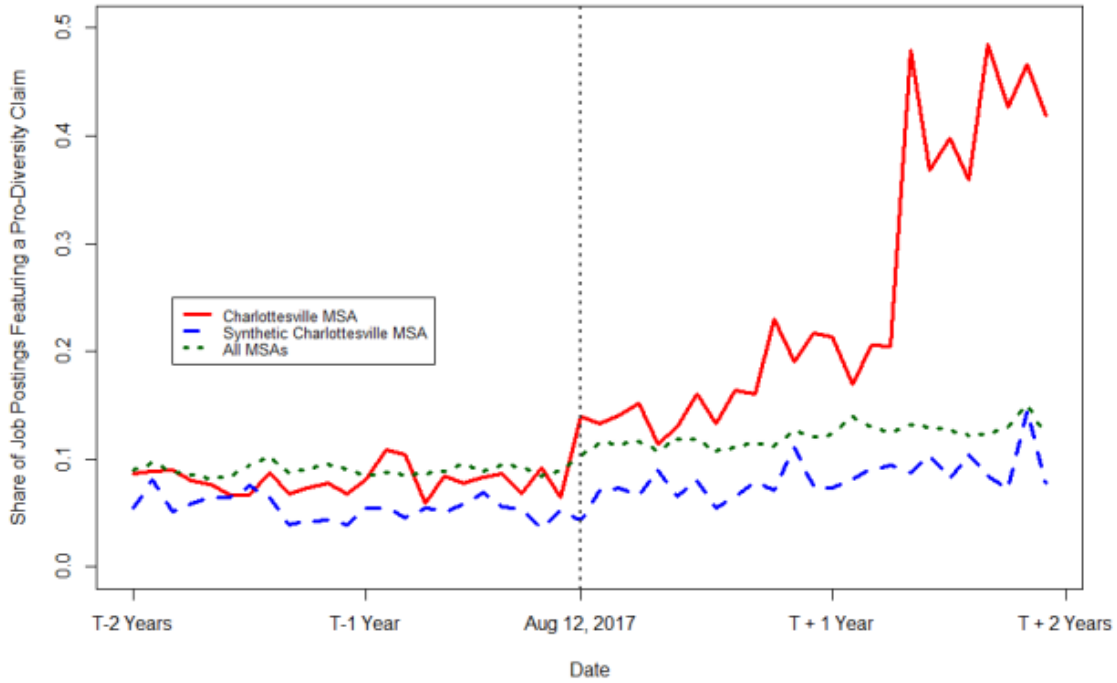


Figure 2.11.10 Excluding postings by UVA, robustness analysis from Roth (2021) shows event-style estimates relative to the hypothetical non-zero pre-trend detectable at the 80% power level.



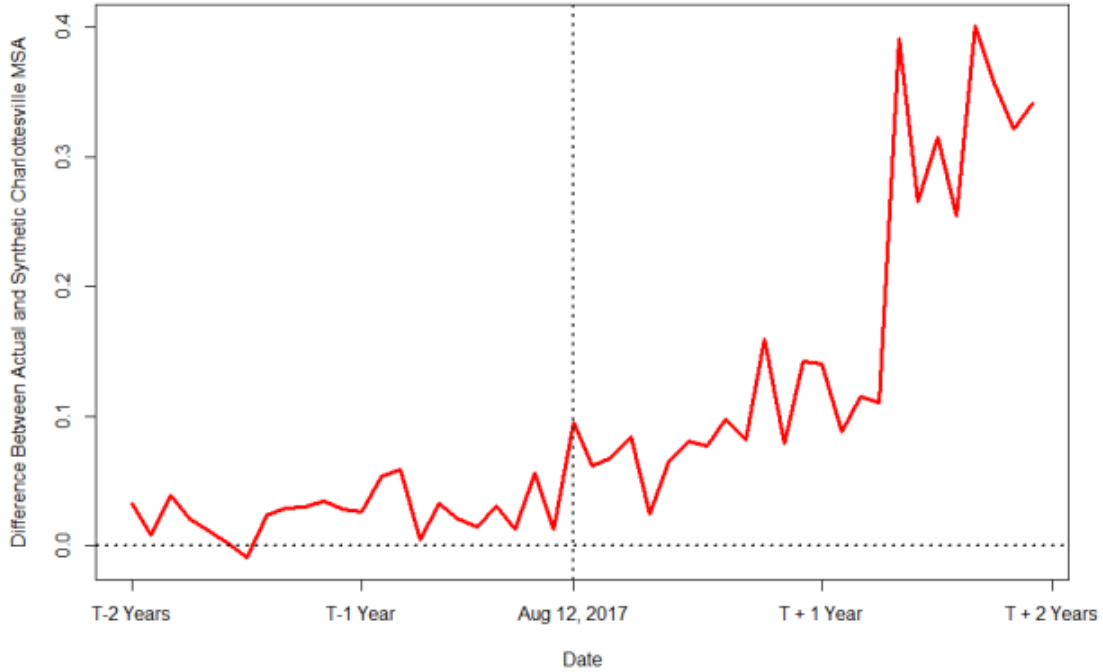
*NOTES:* The circular points and associated confidence intervals are the same time-specific difference-in-differences estimates presented in Figure A6. The triangular points and accompanying trend line indicates the hypothetical non-parallel pre-trend that would be detected at a power level of 80%. This is generated using the methods and software presented in Roth (2021). This line does not cross the 95% confidence intervals of the positive effects estimated in the post-treatment period ( $\delta_2$ ,  $\delta_3$ , and  $\delta_4$ , or the final three points on the right), which is consistent with the conclusion that these estimates are robust to a violation of parallel pre-trends of this size.

Figure 2.11.11 The difference in the share of job postings that included pro-diversity claims in the Charlottesville MSA versus the synthetic Charlottesville MSA grew significantly following the rally.




*NOTES:* Points indicate the monthly share of job postings featuring pro-diversity claims. The share for the synthetic Charlottesville MSA is calculated using the synthetic control method presented in Abadie et al. (2010). The region to the right of the vertical dashed line indicates the post-rally period.

Figure 2.11.12 The difference in the share of job postings that included pro-diversity claims in the Charlottesville MSA versus the synthetic Charlottesville MSA grew significantly following the rally.



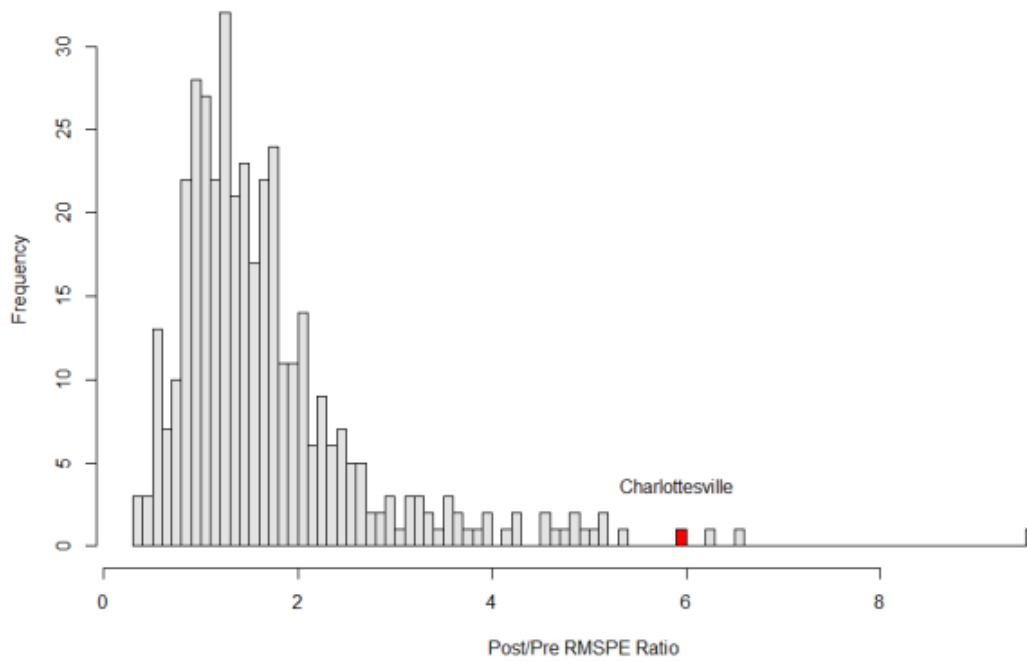
*NOTES:* Points indicate the difference, by month, in the use rates of pro-diversity claims for job postings in the Charlottesville area versus the synthetic control area. The region to the right of the vertical dashed line indicates the post-rally period. The dramatic increase about six months before the close of the study period reflects a significant increase in the use of pro-diversity claims by the University of Virginia that began at the start of 2019.

Figure 2.11.13 These lines illustrate the difference in the share of job postings featuring pro-diversity claims in the actual versus synthetic geographies for all MSAs where RMSPE equal or less than Charlottesville. Charlottesville shown in red.



*NOTES:* This charts the difference in use rates of pro-diversity claims between the actual and synthetic MSA for MSAs in the pool of untreated MSAs (the blue dashed lines), as well as Charlottesville (the bold, solid line marked in red). Consistent with recommendations from Abadie et al. (2010), to increase the readability of the graph, I include only those MSAs with pre-treatment RMSPEs that were less than or equal to that of Charlottesville. Again, this demonstrates that it is very unlikely that the effects detected in Charlottesville, presented in Figure A12, were due to chance.

Figure 2.11.14 The distribution of the ratio of pre- versus post-treatment root mean squared prediction error for all MSAs demonstrates it is very unlikely the observed effect is due to chance.



*NOTES:* This histogram shows the distribution of the the ratio of pre- versus post-treatment root mean squared prediction error for Charlottesville as well as the universe of placebo MSAs. This is the test statistic for the size of the effect of the Unite the Right rally where a larger value indicates a larger effect. The value of this statistics for Charlottesville, labelled and darkened on the far right, is exceptionally large, corresponding to an exact p-value of 0.010, meaning it is very unlikely that the treatment effect observed for Charlottesville, visible in Figure A12, was simply due to chance.

### 2.11.1 Methodology for Recruiting Interviewees

I identified interviewees through the alumni relations office of two universities in the United States. Both are consistently ranked in the top 50 of the US News and World Report college rankings. I undertook this recruiting strategy with the idea that these alumni, given the elite nature of the universities, were likely to hold managerial or other decision-making positions within their employers and were thus likely to have insight regarding if or how their employers adjusted recruiting practices in response to the Unite the Right rally. I identified and contacted all alumni whose current addresses, as per information available through the alumni relations offices, indicated they lived in Charlottesville. I also took a snowball approach where I asked the alumni that replied to my requests if they had colleagues or friends in the community that might be willing to chat with me. For those that were disposed to meet, I conducted the interviews via Zoom or in-person during a research trip I took to Charlottesville in the summer of 2022. In all, I was able to conduct 21 interviews.

### 2.11.2 Methodology for Identifying Pro-Diversity Claims

As mentioned in the body of paper, I code a job posting as featuring a pro-diversity claim, if it meets any one of the following three criteria:

1. It includes the word “diversity,” but not only within the words “diversityjobboard,” “diversityjobs,” “diversitynursing,” or “pharmadiversity.”

2. It includes the word “inclusion” but not the word “inclusion” followed by “in”, “of”, “onto” or “with.”

3. It includes the word “inclusive” but not the phrase “all inclusive” or “inclusive of.”

Here, I lay out the process whereby I arrived at (Steps 1-8) and validated (Step 9) this strategy. This exercise was completed with the full Burning Glass data for the years 2015-2019.

1. Preliminary analysis: My first step in identifying pro-diversity claims was to read through a random sample of 500 job postings. In this preliminary, exploratory exercise I noted that employers made pro-diversity claims using the language of diversity, equity, and inclusion. I noted that pro-diversity claims seemed to most commonly feature the word “diversity,” with related language commonly appearing along with the word diversity, or, in a small number of cases, on its own. Informed by this preliminary analysis, I then systematically examined employers’ use of the nouns “diversity,” “equity,” and “inclusion” and their corresponding adjectives “diverse,” “equitable,” and “inclusive.” As detailed in steps 2-7 below, I first verified that the word “diversity” accurately captured pro-diversity claims and then checked the extent to which job postings that did not feature the word diversity but did feature these other words featured pro-diversity claims.

2. Analyzing “diversity”: I began by analyzing employers’ use of the word “diversity,” which appears in 9.81% of job postings. I took a random sample of 500 job postings that

included the word diversity and, based on posting-by-posting manual analysis of all 500 of these postings, calculated that 93.98% of these were true-positives in the sense that they in fact contained prodiversity claims. Frequently, these involved an employer's assertion that it was "dedicated to," "committed to," "celebrates," "values," "believes in," "thrives on," or "embraces" diversity. I noticed in a small number of cases (<1%) that the only mention of diversity was in relation to the name of the relevant job board, including "diversityjobboard," "diversityjobs," "diversitynursing," or "pharmadiversity." Accordingly, I did not code a job postings as featuring a pro-diversity claim if the word "diversity" only appeared in the context of one of these job board names.

3. Analyzing "equity": I next examined employers' use of the word "equity," which appeared in 1.21% of job postings and 1.00% of jobs that did not feature the word "diversity." I manually checked a random sample of 500 job postings that included the word equity but did not include the word diversity. Again based on on posting-by-posting manual analysis of all 500 of these postings, I calculated that only 9.16% were true positives. For the vast majority of false positives, the word equity referred to financial equity. Unfortunately, there was no common pattern among these false positives which allowed me to systematically identify them. Given this very high rate of false positives, at this step I did not code job postings that included the word equity but did not include the word diversity as containing pro-diversity claims.

4. Analyzing "inclusion": I next examined employers' use of the word "inclusion," which appeared in 2.42% percent of job postings and fewer than 0.46% of job postings that did not include the word "diversity." I manually checked a random sample of 500 job postings that included the word "inclusion" but did not include the word "diversity." Again based on posting-by-posting manual analysis of all 500 of these postings, I calculated that 46% were true



positives. In this case, however, 72.22% of false-positives were easily identifiable as those in which the word “inclusion” was followed by one of the prepositions “in,” “of,” “onto,” or “with.” Of the remaining job postings from this sample of 500, 74.75% were true positives. To double check the efficacy of this approach, I take a fresh random sample of 100 job postings that feature “inclusion,” but were not followed by “in,” “of,” “onto,” or “with” and did not include “diversity.” Consistent with my previous estimate, 76% of these are true positives. Given this high rate of true positives, I coded job postings that featured “inclusion” but did not feature inclusion followed by “in,” “of,” “on,” “onto,” or “with” as featuring pro-diversity claims. I also estimate that this captures about 20% of the true-positives featuring the word equity inasmuch as it detects job postings that featured some combination of the words “equity” and “inclusion” but not the word diversity.

5. Analyzing “diverse”: I next examined the word “diverse,” which appeared in 12.98% of job postings and 11.62% of job postings that did not include the word “diversity.” I manually checked a random sample of 500 job postings that included the word “diverse” but did not include the word “diversity.” Again based on posting-by-posting manual analysis of all 500 of these postings, I calculated that 44% were true positives. In this case, however, there was no systematic way to identify false positives. The issue here was that in many cases “diverse” explicitly characterized the diversity of something that was not demographic diversity, such as distributors, product offerings, or job requirements. This problem also existed to a small extent with “diversity,” but was much more prevalent here. I also noticed at this point, however, that over 30% of the true positives included the word “inclusive,” a detail I return to in Step 7, below. Given this high rate of false positives, at this point I did not code job postings that included the word “diverse” but did not include the word “diversity” as including pro-diversity claims.

6. Analyzing “equitable”: I next examined the word “equitable,” but since this word was extremely rare (appearing in only 0.28% of job postings and fewer than 0.22% of job postings that did not include the word “diversity”), I did not manually check for false positives or otherwise use it to identify pro-diversity claims.

7. Analyzing “inclusive”: Finally, I examined the word “inclusive,” which appeared in 4.82% of job postings and 3.79% percent of job postings that did not include the word diversity. I manually checked a random sample of 500 job postings that included the word “inclusive” but did not include the word “diversity.” Again based on posting-by-posting manual analysis of all 500 of these postings, I calculated that 59.64% were true positives. These true positives very frequently referenced the employer’s inclusive “environment” or “workplace.” In this case, 97.03% of false positives were systematically identifiable as those that include the phrase “all inclusive,” which was used to indicate that a list of applicant requirements or job responsibilities was not exhaustive, and “inclusive of,” which preceded a list of applicant requirements or job responsibilities. After removing these false positives, the true positive rate rises to 97%. To double check the efficacy of this approach, I took a fresh random sample of 100 job postings that feature “inclusive,” but did not include the word “diversity” or the phrases “all inclusive” or “inclusive of” and find that 98% of these were true positives. Accordingly, I code all such job postings as featuring pro-diversity claims.

It is also important to note that, as mentioned in Step 5, by capturing these true-positive uses of “inclusive” I capture a significant number of the true-positive cases that include the word “diverse.” In fact, I estimate that by including these postings, the true positive rate for the remaining job postings that feature “diverse” but not “diversity” falls even further from 44% to

35%. This provides further justification for my decision in Step 5 to not code job postings that feature “diverse” but not “diversity” as including pro-diversity claims.

8. Considering additional diversity-related words: Finally, I examined a range of other diversity-related words, but found that these were very rare. These include “belonging” (appearing in fewer than 0.01% of all job postings), “bipoc” (appearing in 0.01% of all job postings and fewer than 0.01% that did not include the word “diversity”), “black lives matter” (appearing in fewer than 0.001% of all job postings), “dei” (appearing in fewer than 0.001% of all job postings), and “social justice” (appearing in 0.01% of all job postings but far fewer than 0.01% of postings that did not include the word “diversity”). 9. Validating Final Measure: As a final check, I estimated the rate of false positives and false negatives. To check for false positives, I took a random sample of 500 job postings that I had coded as featuring pro-diversity claims. I examined these manually and found that 96% of these were true positives. To check for false negatives, I took a random sample of 500 job postings that I had coded as not featuring pro-diversity claims. I examined these manually and found that 97% of these were true negatives.

### **Chapter 3 Workplace Backlash? Workforce Diversity, Status Threat, and the Contractionary Effects of Pro-Diversity Claims**

#### **Abstract**

Prior research indicates that pro-diversity claims are more effective than non-discrimination claims at expanding and diversifying job applicant pools by appealing to job seekers from marginalized racial groups. Drawing on theories of anti-diversity backlash and group status threat, this paper proposes that growing labor market diversity (i.e., increases in the share of the work- force from historically marginalized groups) may instead cause pro-diversity claims to have a contractionary effect on applicant pools by creating aversion among workers from historically dominant racial groups towards these claims. Results from a pre-registered survey experiment of 2,879 white Americans and exploratory archival analysis show that this “workplace backlash” effect only emerges for non-bachelor’s-holding whites. These findings motivate the development of a general framework for understanding conditions under which pro-diversity claims generate diversified applicant pools that are relatively larger or smaller.

### **3.1 Introduction**

By attracting larger and more diverse applicant pools, employers improve their ability to identify workers that fit their needs (Fernandez, Castilla and Moore 2000). With this objective, firms frequently tailor the language of their recruiting efforts with the hope of expanding applicant pool size, while also diversifying applicant pool composition by increasing representation across social groups (Kalev, Dobbin and Kelly 2006, Kang et al. 2016, Abraham and Burbano 2021). Increasingly, scholars provide insight into how employers might enrich applicant pools by attracting workers from marginalized, rather than dominant, racial groups who, for multiple reasons (Gorbatai, Younkin and Burtch 2021), tend to be severely underrepresented in many labor market contexts (Lin, Aragao and Cobb 2020, Wilson, Miller and Kassa 2021). Broadly speaking, the two contrasting approaches whereby firms seek to convey their receptiveness to applicants from marginalized racial groups can be classified as either pro-diversity claims, which emphasize diversity by celebrating differences, or non-discrimination claims, which emphasize equality by downplaying differences (Apfelbaum, Stephens and Reagans 2016, Dover, Major and Kaiser 2016).

Existing research indicates that, compared to non-discrimination claims, pro-diversity claims more effectively expand and diversify applicant pools. Not only are members of marginalized racial groups more inclined to apply to job postings featuring pro-diversity rather than non-discrimination claims, but members of dominant groups are neither more nor less inclined to apply to postings that make pro-diversity versus non-discrimination claims (Leibbrandt and List 2018, Flory et al. 2021). Even lab experiments showing that members of dominant groups may feel personally “threatened” by pro-diversity claims, meaning they perceive greater discrimination against dominant racial groups or experience adverse physiological reactions when

exposed to these claims, find no evidence that these workers are less attracted to firms that make these claims (Dover, Major and Kaiser 2016, p.61). These studies indicate that employers' objectives for larger and more diverse applicant pools are not incompatible and are both optimized by making pro-diversity claims. They imply, moreover, that as workforces become more diverse, meaning as members of historically marginalized racial groups come to comprise a larger share of the workforce, these size and compositional advantages should increase: In more diverse labor markets, pro-diversity claims will attract more applicants from this growing share of workers from historically marginalized groups, but will not attract fewer applicants from the shrinking, but still significant, share of workers from historically dominant groups. Prior research thus delivers an intuitive conclusion that, in the context of growing workforce diversity, pro-diversity claims, compared to non-discrimination claims, have ever-larger diversifying and expansionary advantages.

I challenge this conclusion by proposing and testing a novel explanation as to why growing labor market diversity may instead cause pro-diversity claims to have a *contractionary* effect on applicant pools. Specifically, I propose and show empirical findings consistent with the argument that perceptions of growing labor market diversity uniquely reduce the extent to which members of dominant groups are attracted to employers that make pro-diversity claims. In building this argument, I draw on theories of anti-diversity “backlash” (Abrajano and Hajnal 2017), which illustrate that, among members of dominant racial groups, perceptions of growing diversity generate feelings of group status threat (Outten et al. 2012), exclusionary attitudes (Abascal 2020), and a consequent reduction in support for parties, policies, and politicians perceived as benefiting marginalized groups (Craig and Richeson 2014*b*, Major, Blodorn and Blascovich 2018, Wetts and Willer 2018). Abrajano and Hajnal (2017) use backlash arguments to challenge conventional wisdom that growing diversity within an electorate has an unambiguously expansionary effect on pro-diversity

political parties. Instead, they demonstrate that growing diversity can have a contractionary effect as anti-diversity backlash causes voters from dominant groups to eschew these parties in favor of those they see as buttressing their historically privileged position atop the social hierarchy. This logic, in turn, has been used to explain the rise of reactionary political movements, which, in this telling, arise to meet this demand for policies and politicians that buttress the at-risk status quo.

I propose that an analogous process, with similarly critical theoretical implications, emerges in labor market contexts as growing workforce diversity reduces the extent to which workers from historically dominant groups are attracted to employers that make pro-diversity claims. Just as prior research shows that feelings of group status threat reduce attraction to parties seen as benefiting marginalized groups, I submit that these feelings spill into labor market contexts in the form of reduced attraction to employers that make pro-diversity claims, which are similarly seen as favoring marginalized groups (Dover, Major and Kaiser 2016). I substantiate this claim with reference to growing evidence that the behavioral implications of workers' sociopolitical preferences are not limited to the political sphere, but also affect where workers decide to supply their labor (Burbano 2016, Bermiss and McDonald 2018, McConnell et al. 2018, Burbano 2021). I propose, moreover, that this backlash effect will not emerge vis-a-vis non-discrimination claims. This is because in the context of waning group status, members of dominant groups begin to see themselves, rather than members of marginalized groups, as victims of discrimination, a phenomenon sometimes called "reverse racism" (Wilkins and Kaiser 2014, Jardina 2019, p.144). In this context, non-discrimination claims, compared to pro-diversity claims, are seen as buttressing the at-risk group status. Thus, perceptions of growing labor market diversity generate a newfound preference for non-discrimination over pro-diversity claims, or, stated differently, an emerging state of the world in which pro-diversity claims, relatively speaking, repel applicants from dominant

social groups. Drawing boundary conditions, I emphasize that this workplace backlash is unlikely to emerge uniformly across a dominant group. Instead, it will be most pronounced among working class members of the dominant group since these workers are more disposed to feel threatened by the perception that the group's historically dominant position is at risk (Du Bois 1935, Manstead 2018, Jardina 2019, p.117). I note that although these workers have rarely been the focus of scholars of human capital strategy, they are frequently critical to firm success.

I examine these predictions in the context of the increasingly racially diverse United States labor market. Whereas white workers, the country's historically dominant racial group, comprised over 80% of the U.S. workforce in 1980, this share is falling rapidly and is expected to sink below 50% by the middle of this century (Toossi 2002). Leveraging this context of growing labor market diversity, I conducted a pre-registered survey experiment on a sample of approximately 3,000 white Americans that was representative of this population in terms of education, political partisanship, and gender. This sample included both women and men. The pre-analysis plan is included as a supporting document and available at this URL: <https://tinyurl.com/4sm8xy79>. I employed a three-by-two design in which subjects indicated their attraction to a job posting that featured either a (1) pro-diversity claim, (2) a non-discrimination claim, or (3) neither of these claims after they had been exposed to information illustrating either (1) the continuity of a white majority workforce or (2) the emergence of a workforce in which workers of color constitute a majority. Consistent with past research, I operationalized working class status in terms of whether subjects did not hold a bachelor's degree (Stephens, Markus and Phillips 2014, p.613-614).

Consistent with my theorized group status threat mechanism, I find that information illustrating growing labor market diversity significantly decreased attraction to the employer making a pro-diversity claim, but only among these non-bachelor's-holding subjects. For these subjects, this



information generated a significant preference for the non-discrimination firm over the pro-diversity firm that did not exist among subjects shown information illustrating the continuity of a white majority workforce. Establishing the external validity of these findings, I present exploratory analysis showing that non-bachelor's-holding subjects, but not bachelor's-holding subjects, residing in local labor markets that experienced greater increases in labor market diversity between 2010 and 2019 were significantly less attracted to the pro-diversity firm. These findings, as well as additional robustness analysis, are consistent with the conclusion that, in the United States, growing labor market diversity is in fact decreasing white working-class job seekers' attraction to employers that make pro-diversity claims.

To demonstrate the theoretical implications of these empirical findings, I develop a framework for understanding conditions under which this "workplace backlash" implies predictions that contrast those implied by prior research. I show that, when accounting for workplace backlash, pro-diversity claims, compared to non-discrimination claims, no longer unambiguously expand applicant pool size. Instead, they often reduce applicant pool size, especially when dominant groups continue to comprise a relatively large share of the workforce. I show, moreover, that the backlash effect accentuates the diversifying effects of pro-diversity claims. Thus, in contrast with delivered research which has implied that employers' expansionary and diversifying objectives are not incompatible, I demonstrate a previously undocumented trade-off in which, as labor markets become more diverse, pro-diversity claims often have a contractionary effect on the size of the applicant pool, but an enhanced diversifying effect on its composition.

I conclude by discussing additional implications and contributions. First, I discuss the possibility that, despite their contractionary effects, pro-diversity claims may benefit firms by filtering out workers from dominant groups who, due to feelings of group status threat, are unlikely to work

well in the diverse workplaces that will result from growing labor market diversity. Avoiding these “misfit” employees plausibly directly reduces turnover (Chatman 1991, Carnahan, Kryscynski and Olson 2017, Bermiss and McDonald 2018), and may indirectly reduce turnover among workers from marginalized groups by reducing the risk that they will experience workplace hostility. Second, I discuss how this paper expands literature regarding the effects of non-pecuniary human capital strategy. Whereas a growing number of studies examine *micro* level contingencies of these strategies, including at the individual (Sauermann and Roach 2014, Carnahan, Kryscynski and Olson 2017) or organizational (Abraham and Burbano 2021) level, this paper demonstrates a process whereby the *macro*-social environment shapes workers’ non-pecuniary preferences, and, by implication, the effects of corresponding non-pecuniary human capital strategy. Finally, this paper advances broader literature regarding anti-diversity backlash among dominant groups. It provides the first evidence, to my knowledge, that anti-diversity backlash extends beyond the political sphere into the labor market, and suggests a novel mechanism whereby recruiting strategy may generate labor market sorting not only by race, but also by racial attitudes.

### **3.2 Theory**

In this section, I first outline prior research, which suggests that as labor markets grow more diverse, pro-diversity claims attract applicant pools that are evermore larger and more diverse than those attracted by non-discrimination claims. I then summarize theories of anti-diversity backlash, and, drawing on these arguments, propose that, among members of dominant groups, perceptions of growing workforce diversity generate a newfound aversion towards employers that make pro-diversity claims. I then discuss why these effects are unlikely to emerge vis-a-vis non-discrimination claims. Finally, I explain why these effects will be most pronounced among working class applicants, since these workers are more reliant on the dominant group identity when forming

self-concept and are thus more disposed to feel threatened by the perception that the group's privileged status is under threat.

### 3.2.1 Prior Research on the Applicant Pool Effects of Pro-Diversity Claims

Pro-diversity claims constitute an increasingly common strategy whereby firms attempt to expand the size and diversify the racial composition of their applicant pools. The defining feature of these claims is a commitment to recognize or celebrate worker differences, especially those based on race, gender, and sexual orientation (Dover, Major and Kaiser 2016). Recent examples include job postings by Elwood Staffing, an Idaho-based event company, claiming to “celebrate diversity at all levels of the organization” and Sony, which, in recruiting for a Nebraska-based retail sales supervisor position, claimed to “embrace diversity.” These pro-diversity claims contrast with non-discrimination claims, which instead avow to downplay or ignore differences by treating applicants equally. Recent examples include a job posting from the management consultancy Accenture claiming, “that no one should be discriminated against because of their differences,” and claims of “equal opportunity” that have proliferated under the Equal Opportunity Act (Dobbin 2009).

Recent field experiments demonstrate that, compared to non-discrimination claims, pro-diversity claims are disproportionately attractive to job seekers from marginalized racial groups. In an experiment conducted in partnership with a large financial services firm, Flory et al. (2021) found that job postings including pro-diversity claims, compared to postings without these claims, more than doubled interest among job seekers from marginalized racial groups. In contrast, A similarly realistic experiment conducted by Leibbrandt and List (2018), found that job postings including non-discrimination claims significantly reduced the number of applicants from marginalized racial groups. These results echo longstanding social psychological theories that members of marginalized groups tend to fare better within and thus prefer environments that

positively emphasize, rather than downplay, group differences (Richeson and Nussbaum 2004, Purdie-Vaughns et al. 2008, Plaut et al. 2018, Ramarajan and Reid 2020).

Existing research, however, has found that members of dominant racial groups are neither more nor less inclined to apply to firms that make pro-diversity versus non-discrimination claims. None of the field experiments cited above, for example, found that dominant group workers were significantly attracted to or deterred by either pro-diversity or non-discrimination claims. Although some work has found that members of dominant racial groups view pro-diversity claims as more exclusionary than non-discrimination claims (Plaut et al. 2011, Dover, Major and Kaiser 2016), this work has not found these workers to be less attracted to employers that make pro-diversity claims. Even work by Dover, Major and Kaiser (2016), which finds that members of dominant racial groups exhibit negative physiological reactions to pro-diversity claims compared to “neutral” claims, found these workers were not less attracted to employers that made pro-diversity claims (see page 61). Older work suggests that workers from dominant groups might even prefer firms that make pro-diversity claims, inasmuch as they associate these claims with a more egalitarian work environment (Williams and Bauer 1994). Delivered research thus suggests that firms’ objectives to both expand and diversify their applicant pools are not incompatible and are both optimized by making pro-diversity claims. It has implied an intuitive conclusion that, as labor markets diversify, these expansionary and diversifying advantages will amplify, because pro-diversity claims will prove exceptionally attractive to the growing share of workers from marginalized groups without reducing attraction among the shrinking but still significant share of workers from historically dominant social groups.

### 3.2.2 Anti-Diversity Backlash

Challenging this conclusion, I propose that growing labor market diversity may cause pro-

diversity claims to instead have a contractionary effect by reducing the extent to which job seekers from historically dominant groups are attracted to these claims. In doing so, I draw on research demonstrating that growing diversity causes members of dominant groups to resist, rather than accommodate or embrace, marginalized groups as well as the politicians, parties, and policies that benefit them. These racial “backlash” theories (Abrajano and Hajnal 2017) originated with claims that racial prejudice on the part of white Americans resulted from fear that black Americans threatened their privileged position atop the racial hierarchy (Du Bois 1935, Blumer 1958, Blalock 1967). These ideas were generalized and refined within intergroup emotion theory (IET), which drew on social identity theory (Tajfel 1982) to argue that group identities provide frames through which individuals observe and respond to social events. According to IET, identification with a group causes individuals to interpret events in terms of their group-level rather than individual-level implications (Maitner, Smith and MacKie 2017). In this telling, members of dominant groups experience positive self-image from their position atop the social hierarchy, and their resentment among dominant groups toward marginalized groups arises from feelings that these groups threaten this dominant status.

Perceptions of relative group sizes are a critical driver of these feelings of racial status threat. Studies in many contexts show that dominant racial group animosity towards marginalized racial groups increases as the relative size of these marginalized groups increases (Taylor 1998, Quillian 1995, Craig and Richeson 2014*a*). These feelings in turn shape political preferences, pushing voters from dominant groups to support policies and politicians they believe will perpetuate their racial group’s privileged position, and to resist those they perceive as raising the status of marginalized groups (Craig and Richeson 2014*b*, Wetts and Willer 2018, Jardina 2019). As mentioned above, political scientists have drawn on these findings to challenge conventional wisdom that growing

diversity expands the voting base of pro-diversity parties, instead showing growing racial diversity can have a contractionary effect by causing voters from dominant groups to eschew these parties (Abrajano and Hajnal 2017).

### 3.2.3 Workplace Backlash?

I argue that an analogous process plays out in labor markets where perceptions of growing labor market diversity generate newfound aversion among workers from dominant groups towards firms that make pro-diversity claims. These claims seek to draw contrast with companies without these claims where worker diversity is presumably muffled under a dominant-group-favoring, historical status quo. In the absence of threat, workers from dominant racial groups are indifferent towards, and perhaps even attracted by (Williams and Bauer 1994), the idea of working for a firm which espouses a pro-diversity claim. But as perceptions of growing labor market diversity create anxiety that group status is under threat, workers experience newfound dissonance with respect to working for and thus contributing to the success of a firm that celebrates the very diversity they see as undermining their now-at-risk group status. This proposed process is consistent with growing evidence that sociopolitical preferences significantly shape where workers decide to supply their labor (Bermiss and McDonald 2018, McConnell et al. 2018, Burbano 2021). If perceptions of growing diversity generate anti-diversity sociopolitical preferences, these preferences will manifest in labor market contexts in the form of newfound aversion to employers that make pro-diversity claims.

This newfound aversion, however, is unlikely to emerge with respect to non-discrimination claims. This is because in the context of growing diversity, members of dominant groups increasingly see themselves, rather than members of historically marginalized groups, as victims of discrimination (Wilkins and Kaiser 2014, Jardina 2019, p.144). Multiple surveys conducted in the

United States over the past decade, for example, demonstrate that a growing number of whites, the country's dominant racial group, believe that discrimination against whites is as severe as discrimination against marginalized racial groups (Cox, Galston and Jones 2011, Gonyea 2017). Theoretical explanations for these perceptions of "reverse discrimination" draw on loss aversion, or the idea that the psychological impact of lost status for members of dominant racial groups is especially severe due to their group's historical position atop the social hierarchy (Norton and Sommers 2011). In this context, the meaning of non-discrimination, which previously may have been understood simply as promising non-discrimination against marginalized groups (Dobbin 2009), comes to be seen as promising non-discrimination against dominant groups. Non-discrimination claims, moreover, transmit, compared to pro-diversity claims, a relatively colorblind, rather than multicultural, ideology, which members of dominant groups are more likely to view as buttressing, rather than undermining, their dominant position in the social hierarchy (Knowles et al. 2009). Thus, I propose that as workers from historically marginalized groups come to comprise a larger share of the workforce, workers from historically dominant groups become relatively deterred by pro-diversity compared to non-discrimination claims.

### 3.2.4 Social Class as a Key Boundary Condition

It is unlikely, however, that this workplace backlash will emerge uniformly across workers from dominant groups. As explained above, racial backlash arguments build from the assumption that dominant group members' attachment to their shared identity arises from the fact that they derive positive self-image from its position atop the social hierarchy. It follows that the perceived erosion of this dominant position, and by implication the erosion of this source of positive across-group comparison, will be felt most acutely among members of dominant groups for whom this comparison is especially important in forming positive self-image (Tajfel 1982). For those that rely

disproportionately on other identities to form positive self-image, group status threat and the consequent aversion to pro-diversity claims will be muted or perhaps not emerge at all.

Prior research indicates that susceptibility to group status threat will be especially high among “working class” applicants from dominant racial groups. Although precise definitions vary, working class, compared to middle or upper class, workers enjoy lower wages, less upward mobility, and less formal education (Pellegrin and Coates 1957, Leana, Mittal and Stiehl 2012). Stephens, Markus and Phillips (2014, p.613) define the working class as “individuals in contexts on the bottom half of the social class divide, including people who have attained less than a four-year college degree or who have relatively lower income or lower-status occupations.” Among working class members of historically dominant racial groups, racial identity has been shown to play an especially critical role in forming self-concept, since it compensates for the professional rewards that are largely unavailable due to their disadvantaged location within the broader capitalist structure (Manstead 2018, Myers 2019). This compensating logic was captured by W.E.B. Du Bois who, observing working class whites in the antebellum American South, characterized their attachment to racial identity as a non-pecuniary “psychological wage” that compensated for their precarious economic condition (Du Bois 1935, p.700). Consistent with Du Bois’s observation, Jardina (2019, p.117) concludes that “working class whites have been able to use race as a means to elevate and distance their own group from those at the bottom of the social order,” and are therefore, “especially reactive to threats to the advantages and status they accrue from being white.” In short, if group status threat is the mechanism connecting perceptions of growing labor market diversity to decreased attraction to pro-diversity claims, then this relationship should be strongest among working class members of dominant social groups.

Before progressing to empirical analysis, I underscore the strategic relevance of these



workers. Research in strategic human capital generally focuses on highly skilled workers, such as scientists, gifted managers, or other “stars” who have the skills to provide and/or acquire advantage-giving, firm specific knowledge (Coff 1997, Wang and Barney 2006). Less skilled workers, in contrast, are often assumed to be relatively interchangeable and thus less strategically consequential. It is worth highlighting, however, that, even in geographies with the largest share of skilled workers, the majority of workers are working class. In the United States, for example, around 60% of workers do not hold a bachelor’s degree (Brundage 2017). This rate is similar across Europe (Educational Attainment 2020). Thus, even if workplace backlash is largely limited to the working class, it will permeate large portions of the labor market and affect a large proportion of firms.

The existence of this proposed backlash effect would imply that, in contrast with predictions from delivered theory, growing diversity will not unambiguously enhance the applicant-pool-expanding effects of pro-diversity claims. Instead, it would imply that pro-diversity claims may have a contractionary effect on applicant pools as workers from historically dominant groups, especially workers within these groups that hail from the working class, increasingly refrain from applying to firms that make these claims. In the following three sections (Section 3-6) I present the empirical context and experimental design with which I test for this effect, and empirical results indicating that it exists only among working class whites. In Section 7, I present a simple framework with which I integrate this empirical finding with prior work to explain the conditions under which workplace backlash causes pro-diversity claims to have a contractionary effect on applicant pools.

### **3.3 Empirical Context**

I examine these predictions in the context of the increasingly raciall diverse United States labor market. Whereas workers from the country’s dominant racial group (white workers) comprised

over 80% of the U.S. workforce in 1980, the share of workers from historically marginalized racial groups (workers of color) is rising rapidly and is expected to surpass 50% by the middle of this century (Toossi 2002). This labor market transformation is part of broader demographic transformation in which 2010-2019 marked the first decade on record in which the number of white Americans fell (Tavernise and Gebeloff 2021). This broader demographic transformation has been shown to generate feelings of group status threat among white Americans (Outten et al. 2012). Importantly these feelings have been shown to be driven by concerns of waning group status, rather than individual material self-interest caused by economic competition (Mutz 2018, Margalit 2019).

These feelings are especially acute among working class whites. Compared to their upper- and middle-class counterparts, these workers face lower opportunities for career advancement, greater wage stagnation, as well as upward trends in chronic illness, suicide, and substance abuse (Case and Deaton 2015). Surveys show that working class white Americans are significantly more likely than their upper- and middle-class counterparts to feel threatened by growing diversity (Jones et al. 2016). Qualitative fieldwork further confirms these workers' fear of waning status and illustrates their growing aversion towards public policies they perceive as unfairly advantaging marginalized groups (Cramer 2016, Hochschild 2016, Gest 2016). In these accounts, working class whites frequently express frustration that they, rather than members of marginalized racial minority groups, are the true victims of discrimination. Thus, in the context of the United States labor market, workplace backlash effect should be most pronounced among this group.

### **3.4 Survey Experiment**

I test for the existence of my proposed backlash effect using a two-dimensional survey experiment in which I randomize exposure to information illustrating varying trends in labor market diversity and

then measure attraction to randomized job posting language that makes either a pro-diversity claim, a non-discrimination claim, or neither of these claims. This approach to manipulating perceptions of demographic change has gained popularity as a preferable alternative to archival analyses of demographic change, which are less equipped to provide credible grounds for causal inference (Wetts and Willer 2018, Abascal 2020).

#### 3.4.1 Target Sample

I directed the survey company Qualtrics to recruit a sample of 3,000 non-Hispanic white Americans that was representative of this group in terms of educational attainment, political partisanship, and gender. Based on estimates from the U.S. Census Bureau, I stipulated that 40% of subjects hold at least bachelor's degree (Educational Attainment 2019). Based on estimates from the Pew Research Center, I stipulated that the sample feature 30% Democrats, 35% Republicans, and 35% Independents (Wide 2018). This sample included both men (50%) and women (50%).

#### 3.4.2 Experimental Procedure

To conduct the experiment, I posed as a human resources consultancy (hereafter Firm A) and instructed subjects to provide feedback on recruitment material. To increase the realism of the subject's experience, I changed the Qualtrics survey link URL to reflect this consultancy's name and instructed Qualtrics to ensure that this name was used in their proprietary subject recruitment efforts. When subjects began the survey, they were greeted with the company name, given a single line describing the company as a human resource consultancy, then answered questions measuring gender, education level, age, political partisanship, race, and Hispanic identity. These were used by Qualtrics to fulfill my requested demographic quotas. Subjects were then informed they would be asked to perform three tasks on behalf of the company. I used these tasks to administer treatments and measure outcomes of interest.

I manipulated perceptions of workforce diversity via the first task, in which I instructed subjects to review two pages from an ostensible newsletter prepared by Firm A for its clients. The first page was a table of contents, identical for all subjects. Subjects were then randomized to receive a second page of content illustrating either that white workers will continue to be a majority of the workforce or that white workers will become a minority. Similar to the intervention used by Abascal (2020), this content featured real demographic projections, in this case workforce demographic projections from the United States Department of Labor Statistics (Toossi 2002), and manipulated perceptions of workforce diversity by subtly adjusting the date range (x-axis) of the graphic in each condition. In the white majority condition, the graphic illustrated trends between 2005 and 2025, which showed the persistence of a white majority workforce. In the white minority condition, the graph illustrated demographic trends between 1980 and 2070, which showed the disappearance of the white majority (i.e., growing workforce diversity). Figure 1 summarizes these manipulations, while the pre-analysis plan includes the exact manipulations.

Next, subjects provided feedback on language that Company A's ostensible client (hereafter Company B) planned to include in a job posting. Subjects were instructed to imagine themselves as prospective employees and to consider whether they would be interested in working for Company B based on this information. Subjects were randomly assigned to view information featuring a pro-diversity claim, a non-discrimination claim, or generic information that made neither of these claims. This generic information was included in all treatments and was sufficiently vague that the company could be seen as a potential employer for all subjects regardless of education or skill. Table 1 summarizes these treatments. The pro-diversity claim mentioned multiple dimensions of diversity, in this case not only racial, but also gender and culture differences. This better operationalized the broader theoretical construct of a pro-diversity claim and better matched how

such statements appear in actual job postings. This realism limited the extent to which effects might have been driven by perceived strangeness of the claim had it mentioned only one of these dimensions. The non-discrimination claim was nearly identical, but instead claimed that the firm would not discriminate based on these differences. Inasmuch as it did not make the same mention of race, gender, and culture, the generic description did not serve as a clean control, but provided an additional benchmark.

After viewing one of these three descriptions, subjects answered five survey items measuring attraction to Company B. This five-item index was developed by Highhouse, Lievens and Sinar (2003) as a robust, survey-based measure of “employee recruitment and organization choice.” This index has been widely used. The index’s five items include: *1. For me, this company would be a good place to work, 2. I would not be interested in this company except as a last resort, 3. This company is attractive to me as a place of employment, 4. I am interested in learning more about this company, and 5. A job at this company is very appealing to me.* Subjects responded to these questions on a five-point Likert scale ranging from “Strongly Disagree” to “Strongly Agree.” This outcome is less direct than behavioral measures of organizational attraction, such as reservation wage (Burbano 2016) or demonstrated interest in job ads (Abraham and Burbano 2021, Snellman and Younkin 2021), but is more appropriate for this multidimensional, information treatment design in which the identities of the ostensible provider of information regarding workforce diversity and the ostensible hiring firm needed to be separate. Subjects then answered a set of questions measuring perceptions of the firm’s demographics, as well as perceptions of discrimination within the firm. In a final task, subjects edited pages from an ostensible business overview prepared by Company B. In the “PreAnalysisPlanDeviations.pdf” document I discuss why analyses of performance on this task is not included in the main text.

### 3.4.3 Measures

**Independent Variables:** Figure 2 summarizes the 3x2 treatment design. The variable *ProDiversity<sub>i</sub>* is a dummy variable equal to one if subject *i* received the pro-diversity description of Company B. *Generic<sub>i</sub>* is a dummy variable equal to one if the subject received the generic description of Company B. I treat the non-discrimination condition as the excluded comparison group to facilitate my central theoretical comparison of interest: pro-diversity versus non-discrimination claims. *WhiteMinority<sub>i</sub>* is a dummy variable equal to one for subjects in the white minority condition.

**Dependent Variable:** *Attraction<sub>i</sub>* is a value created by averaging subject *i*'s scores on each item in the organizational attractiveness index. Each response is coded from 0 to 4 (*Strongly Disagree=0, ..., Strongly Agree=4*), such that *Attract<sub>i</sub>* ranges from 0-4. The value for the second question in this index, *I would not be interested...last resort*, is reverse coded. Like past papers using this index, I treat this as a proxy for intention to seek employment at the firm. In this case, for example, if a subject were to indicate a score of 2.7 for the pro-diversity firm, and a score of 2.3 for the non-discrimination firm, I would infer she is more inclined to apply to the pro-diversity firm. A limitation of this measure is that it does not map precisely into probability space. Accordingly, my exposition and discussion of results emphasizes the direction of differences in attraction levels across the claims and how these differences change upon exposure to the white minority information.

**Working Class:** Consistent with Stephens, Markus and Phillips (2014), I operationalize working class with whether the subject holds at least a bachelor's degree.

## 3.5 Results

### 3.5.1 Summary Statistics and Balance Tests

Table A1 in the appendix presents summary statistics and balance tests. After Qualtrics removed low quality respondents, the full sample included 2879 subjects. See the appendix for how Qualtrics identified these low-quality subjects. The table shows that the desired education and partisanship

quotas were achieved and reports balance tests which illustrate randomization was completed successfully. Table A2 shows the mean and standard deviation of *Attraction* for the overall sample, 2.17 and 0.84 respectively, as well as for each of the six treatment conditions.

### 3.5.2 Pre-Registered Analysis

**Main Effects:** I estimate the main effects of the company descriptions with

$$Attraction_i = \beta_0 + \beta_1 ProDiversity_i + \beta_2 Generic_i + \epsilon_i \quad (1)$$

where  $\beta_1$  is the average treatment effect of the pro-diversity condition relative to the non-discrimination condition,  $\beta_2$  is the effect of the generic condition relative to the non-discrimination condition,  $\beta_0$  is the intercept, and  $\epsilon$  is the error term. I estimate this and all subsequent models using ordinary least squares and heteroskedasticity-robust standard errors. Column 1 of Table A14 reports the estimate of Equation 1. Compared to subjects in the non-discrimination condition, subjects in the pro-diversity condition reported organizational attraction scores that were on average 0.18 lower ( $p < 0.001$ ), and those in the generic condition reported scores 0.28 lower ( $p < 0.001$ ). This shows that, across the full sample, subjects were most attracted to the non-discrimination firm and more attracted to the pro-diversity firm than the generic firm.

Next, I test for my theorized backlash effect by examining whether the white minority information dampened attraction to the pro-diversity claim. Specifically, I estimate

$$Attraction_i = \beta_0 + \beta_1 ProDiversity_i + \beta_2 Generic_i + \beta_3 WhiteMinority_i + \beta_4 ProDiversity_i * WhiteMinority_i + \beta_5 Generic_i * WhiteMinority_i + \epsilon_i \quad (2)$$

where  $\beta_4$  indicates the relative difference in attraction towards the pro-diversity versus non-discrimination firm in the white minority condition versus the white majority condition and  $\beta_5$  represents the analogous comparison for the non-discrimination versus generic firm. The estimate of Equation 2 is reported in column 3 of Table A14. The white minority information decreased

attraction to the pro-diversity firm relative to the non-discrimination firm by 0.10, but this effect was not statistically significant ( $p \approx 0.18$ ). This information induced only a very small (0.02) decrease in attraction towards to the generic firm. Column 3 of Table A3 shows these results do not change with the inclusion of controls for age, gender, and Republican partisanship. Figure A1 summarizes these patterns visually. Consistent with a backlash effect, there is a visible but small decrease in attraction to the pro-diversity claim, with no change in reaction to the non-discrimination claim, but this reduction in attraction is not statistically significant.

***Subgroup Analysis by Bachelor's-Holding Status:*** Next, I conduct preregistered subgroup analyses by whether subjects hold bachelor's degrees, which, as explained above, operationalizes working class identity. As outlined above, my theory predicts the backlash effect will be especially pronounced among working class subjects. Columns 4 and 5 report separate estimates of Equation 2 for bachelor's-holding and non-bachelor's-holding subjects. Column 4 shows that, among non-bachelor's-holding subjects, the white minority information significantly decreased attraction to the pro-diversity relative to the non-discrimination firm ( $\beta_4 = 0.23, p < 0.05$ ). Column 5, in contrast, reports these estimates for subjects who hold a bachelor's degree and shows no significant effects. I also compare these subgroup effects using a triple-interaction model. This is reported in column 6 and shows that the difference in the effect of the white minority information on non-bachelor's-holding versus bachelor's-holding subjects' relative attraction towards the pro-diversity firm versus the non-discrimination firm is 0.40 standard deviations larger ( $\beta_4=0.34, p < 0.05$ ). Again, I replicate these results including demographic controls as shown in Columns 4-6 of Table A3. Figure 3 summarizes this analysis visually. Consistent with a backlash explanation, the left plot illustrates that exposure to the white minority condition significantly decreased non-bachelor's- holding subjects' attraction towards the pro-diversity firm, while generating a slight, if insignificant, increase in



attraction to the non-discrimination firm. Attraction to the generic firm remained almost totally unchanged. The right plot, in contrast, illustrates that exposure to this information did not induce significant change in bachelor's-holding subjects' attraction to any of the three firms. This analysis reveals that the suggestive, but insignificant results summarized in Figure A1 masked a significant divergence in reactions by education level.

Column 4 illustrates that, among non-bachelor's-holding subjects, this information generated a newfound preference for the non-discrimination claim over the pro-diversity claim. In the white majority condition, non-bachelor's-holding white workers were not significantly more attracted to the non-discrimination over the pro-diversity firm (difference of 0.09,  $p \approx 0.25$ ). Exposure to the white minority information, however, created a clear, statistically significant preference for the non-discrimination firm over the pro-diversity firm (difference of 0.32,  $p < 0.05$ ). Comparison with the generic firm shows a similar preference change. Whereas subjects in the white majority condition had a significant preference for the pro-diversity firm over the generic firm (difference of about 0.19,  $p < 0.05$ ), exposure to the white minority information erased this preference (difference of about 0.01,  $p \approx 0.84$ ). Thus, for non-bachelor's-holding subjects, the white minority information uniquely dampened attraction to the pro-diversity firm, generating a new, significant preference for the non-discrimination firm over the pro-diversity firm while also erasing the significant preference for the pro-diversity firm over the generic firm.

### **3.6 Exploratory Analysis**

#### **3.6.1 Examining Alternative Mechanisms**

I propose that this newfound aversion to pro-diversity claims follows from fear of waning group status. The fact that I detect this effect only within a sub-population within the wider white American population that existing research has illustrated is particularly likely to experience this group status threat is consistent with this mechanism. I do not, however, include survey questions

explicitly measuring perceptions of group status threat. Not only would such questions decrease the realism of my experimental tasks, but it is already well documented by many prior studies that exposing white Americans to information regarding demographic change generates group status threat (Outten et al. 2012, Craig and Richeson 2014 $a,b$ , Major, Blodorn and Blascovich 2018, Wetts and Willer 2018). I do, however, conduct additional tests to address the alternative explanation that the white minority condition changed workers' beliefs regarding unobservable characteristics of firms that make pro-diversity claims, and that these beliefs, rather than fear of diminished group status, reduced attraction to the pro-diversity firm.

***Perceptions of Firm Demographics:*** First, I consider the possibility that the white minority information increased subjects' perception of the relative presence of workers or bosses of color in the pro-diversity versus non-discrimination firm. Such an effect would suggest the possibility that the decrease in attraction toward the pro-diversity firm caused by exposure to the white minority information arose not from a newfound aversion to the pro-diversity claim but was rather an indirect reflection of subjects' preference for co-racial workers and/or bosses. To examine perceptions of racial demographics I create two dummy variables indicating whether or not subjects agree or strongly agree that Company B is *more likely than other companies to employ racial minorities, or be led by a boss who is a racial minority*. For each of these outcomes, I reestimate Equation 1 and then reestimate Equation 2 for the full sample and for the education subgroups. These estimates are presented in Tables A4 and A5. Table A4 shows that subjects were 5 percentage points ( $p < 0.01$ ) more inclined to report that the pro-diversity firm, compared to the non-discrimination firm, was more likely than other companies to employ racial minorities. The white minority information did not significantly affect this perception. Table A5 shows that subjects did not perceive the pro-diversity firm as significantly more or less likely to be led by a boss who was a racial minority, and

that this perception did not vary significantly with exposure to the white minority information. These patterns are consistent with the proposed backlash mechanism, showing that the white minority information did not shift subjects' beliefs regarding the high prevalence of workers of color within these firms, but suggesting that this information generated newfound aversion to working alongside these workers.

Second, I examine perceptions of political partisanship within the firms. This explores the possibility that the decrease in attraction toward the pro-diversity firm caused by exposure to the white minority information arose from a newfound aversion to work alongside or under Democrats, who tend to support policies and politicians deemed more beneficial to marginalized groups. To measure these perceptions, I included a survey question asking, "*If you had to guess, would you say that most employees at [Firm B] are Democrats, Republicans, or Independents?.*" I created a dummy equal to one when subjects answered "Democrats." To measure perceptions of the boss's partisanship, I asked, "*If you had to guess, would say the boss of [Firm B] is a Democrat, Republican, or Independent?.*" I created a dummy variable equal to one when subjects responded "Democrat." Analyses of these variables are shown in Tables A6 - A7. Compared to the non-discrimination firm, the pro-diversity firm was perceived as 5 percentage points ( $p < 0.01$ ) more likely to have a workforce composed mostly of Democrats and 8 percentage points ( $p < 0.01$ ) more likely to be led by a Democrat. As before, these perceptions did not change significantly with exposure to the white minority information. These patterns are again consistent with the proposed backlash mechanism, showing that the white minority information did not shift subjects' beliefs regarding the high prevalence of left-leaning employees within these firms, but suggesting that this information generated newfound aversion to working alongside these workers.

*Perceptions of Anti-White Discrimination:* Next, I consider the possibility that the white minority information increased subjects' beliefs that they would suffer discrimination within the pro-diversity firm. In this alternative telling, the reduced attraction among non-bachelor's-holding workers towards the pro-diversity firm in response to the white minority information did not arise from newfound anxiety regarding group status, but rather newfound anxiety that they would personally experience discrimination as a white applicant or employee. Similar to Dover, Major and Kaiser (2016), I explore perceptions of anti-white discrimination with survey items measuring whether respondents believed Company B was likely to *discriminate against white workers, unfairly favor minority workers*, and whether Company B *seems like a fair place to work*. Each of these items is measured on a 5-point Likert scale and I again create a dummy variable for each indicating agreement or strong agreement. These results are presented in Tables A8 - A10. These show no evidence that subjects perceived different levels of discrimination or unfairness across any of the three firms, and, again, no evidence that the white minority information changed this perception.

### 3.6.2 Examining External Validity with Archival Data

Although the randomized firm descriptions correspond to language that applicants commonly confront in actual job postings, it is less clear to what extent the experimental manipulations of perceived trends in labor market diversity correspond to actually experiencing growing labor market diversity. Accordingly, I examine whether subjects' relative attraction to the pro-diversity firm is related to actual diversity growth in their local labor markets of residence. Using the longitude and latitude of the device on which they completed the survey as a proxy for subjects' location of residence, I identify the Metropolitan or Micropolitan Statistical Areas (MSAs) in which each subject resides. MSAs are geographic designations by the United States Census Bureau and can be thought of as local labor markets within the broader United States economy. Subjects in my sample

came from all fifty states and 455 different MSAs. Using the Census Bureau's American Community Survey five-year estimates of MSA-level racial demographics, I calculate the percentage change in the share of the population in each subject's MSA of residence that identify as people of color, meaning those that do not identify as non-Hispanic white, from 2010 to 2019. By these estimates, subjects in my sample were living in MSAs that experienced, on average, an approximately 14% increase in the share of the local population that identified as people of color between 2010 and 2019. Figure A2 illustrates the distribution of this value across subjects.

I then re-estimate Equation 2, substituting *WhiteMinority<sub>i</sub>* with the percentage change in the share of residents that are people of color. This allows me to estimate the heterogeneous treatment effects of the different firm descriptions with respect to local diversity growth. I again estimate the model across the whole sample and then by bachelor's-holding status. These results are presented in Table A11 and are remarkably consistent with the findings from the experiment. Among bachelor's-holding subjects, a ten percent change in the share of the local population that identify as people of color is associated with an approximately 0.10 decrease in relative attraction to the pro-diversity versus the non-discrimination claim ( $p < 0.05$ ). As in the purely experimental analysis, bachelor's-holding subjects' relative attraction to the pro-diversity versus non-discrimination claim is not significantly related to local demographic change. For robustness, I re-estimate the relationship among non-bachelor's-holding subjects while including individual-level (partisanship, gender, and age), and/or MSA-level (average age, median earnings, percent of residents who hold bachelor's degrees, and population) controls, and/or state fixed effects. I also re-estimate this model where I control for the percentage change in each of these macro-level variables rather than their levels. As shown in Table A12, the relationship is robust to these ten alternative specifications.

Figure 4 illustrates predicted attraction to the three firms across different levels of local di-

versity change for non-bachelor's-holding subjects. Whereas predicted attraction towards the non-discrimination and pro-diversity firms was essentially identical at low levels of diversity change, greater diversity change was associated with a preference bifurcation where subjects came to significantly prefer the non-discrimination firm over the pro-diversity firm. As in previous analysis, this growth also erased subjects' preference for the pro-diversity firm over the generic firm. This operationalization of growing workforce diversity is not exogenous, but the consistency of patterns across this hybrid archival-experimental analysis and the main, purely experimental, analysis support the conclusion that the latter are not artifacts of the way I manipulate perceptions of growing labor market diversity.

### **3.7 Implications for Applicant Pool Size and Composition**

Having provided empirical evidence consistent with my proposed backlash effect, I now illustrate this effect's theoretical implications. Specifically, I present a framework with which I delineate conditions under which this effect causes pro-diversity claims to have a contractionary effect on applicant pools. A key objective of this exercise is to integrate my empirical finding with prior work on pro-diversity claims, which has examined more diverse subject pools (Leibbrandt and List 2018, Flory et al. 2021, Abraham and Burbano 2021).

#### **3.7.1 Basic Framework**

I assume a workforce composed of potential applicants from marginalized racial groups  $m$  who comprise a share  $\alpha \in (0, 1)$  of the total workforce and potential applicants from dominant racial groups  $d$  who comprise a share  $1 - \alpha$  of the workforce. These two types of workers,  $\tau \in \{m, d\}$ , are recruited by firms that are equal in every respect except their decision to make job postings that make either pro-diversity claims  $pd$  or non-discrimination claims  $nd$ ,  $\pi \in \{pd, nd\}$ . Potential applicants observe these claims and apply to a given firm with a probability  $p_{\tau,\pi} \in [0, 1]$ . Based

on existing research discussed above, I assume that potential applicants from marginalized groups are more inclined to apply to postings featuring pro-diversity rather than non-discrimination claims, such that  $p_{m,pd} > p_{m,nd}$ , and more inclined than workers from dominant groups to apply to postings featuring pro-diversity claims, such that  $p_{m,pd} > p_{d,pd}$ . The total size of the applicant pool attracted by the claim  $\pi$  is simply the share of all possible applicants who apply to the job, and can be expressed as the sum of the share of applicants from marginalized groups and the share of applicants from the dominant group. Formally, this can be expressed

$$Size_{\pi} = Size_{m,\pi} + Size_{d,\pi} = p_{m,\pi}(\alpha) + p_{d,\pi}(1 - \alpha).$$

Similarly, the diversity of the applicant pool, meaning the share of the applicant pool from marginalized groups, can be expressed

$$Diversity_{m,\pi} = \frac{Size_{m,\pi}}{Size_{m,\pi} + Size_{d,\pi}} = \frac{p_{m,\pi}(\alpha)}{p_{m,\pi}(\alpha) + p_{d,\pi}(1 - \alpha)}.$$

### 3.7.2 Assuming No Backlash

I first consider a state of the world where there is no workplace backlash, meaning labor market diversity does not reduce the extent to which workers from dominant groups are attracted to firms that make pro-diversity claims. This is the state of the world implied by existing research. Based again on prior research, as well as results from the white majority condition in my experiment, I assume that dominant group workers are indifferent between the pro-diversity and non-discrimination claims. For illustration, I select values for  $p_{\tau,\pi}$  that satisfy the assumptions above and roughly match the magnitude of the attraction rates identified in past field studies ( $p_{m,pd} = 0.40$ ,  $p_{m,nd} = 0.30$ ,  $p_{d,pd} = p_{d,nd} = 0.35$ ). These patterns hold for any values of  $p_{\tau,\pi}$  that satisfy the conditions outlined above. I then calculate the size of the applicant pool attracted by each of the two claims as workers from marginalized groups come to comprise a larger share of the labor market (i.e., for growing levels of

$\alpha$ ). These results are presented in the left panel of Figure 5 and illustrate that, compared to non-discrimination claims, pro-diversity claims lead to unambiguously larger applicant pools and this difference grows as workers from marginalized groups come to constitute a larger share of the labor market. Similarly, the right panel shows that the share of applicants from marginalized groups is always greater in the applicant pool attracted by the pro-diversity claim. In this state of the world, firms' expansion and diversifying objectives are both compatible and are both optimized with the pro-diversity claim.

### 3.7.3 Allowing for Backlash

Next, I demonstrate how this conclusion changes when accounting for the workplace backlash effect identified in my empirical analysis. Specifically, I redefine dominant group workers' attraction to the pro-diversity claim,  $p^*_{d,pd} \in [0, 1]$ , such that this attraction is negatively correlated with the share of workers from marginalized groups,  $\alpha$ . Formally, I re-define dominant group members' propensity to apply to the pro-diversity firm as

$$p^*_{d,pd} = p_{d,pd} - \rho\alpha$$

where  $\rho$  represents the extent of backlash experienced by the dominant group. Using  $p^*_{d,pd}$ , I then re-estimate the size and homogeneity of the applicant pool at three different values of  $\rho$  ( $\rho = 0.10$ ,  $\rho = 0.20$ , and  $\rho = 0.30$ ). I then overlay these curves on those illustrated in Figure 5 above. These patterns are shown in Figure 6 and demonstrate the key theoretical implication of my empirical findings. In contrast to predictions provided by existing theory, allowing for backlash effect implies that growing labor market diversity causes the pro-diversity claim to generate applicant pools that are sometimes smaller than those attracted by the non-discrimination claim. Specifically, this shows that for the different values of  $\rho$ , there is a level of labor market diversity (i.e., a value of  $\alpha$ ) below



which the pro-diversity claim attracts a smaller applicant pool than the non-discrimination claim.

This result can be generalized to show that this threshold can be expressed

$$\alpha^* = 1 - \frac{p_{m,pd} - p_{m,nd}}{\rho},$$

which formalizes the intuition that the pro-diversity claim tends to attract a smaller applicant pool than the non-discrimination claim when the share of workers from the marginalized group is smaller, marginalized group workers' preferences for the pro-diversity over the non-discrimination claim is smaller, and the workplace backlash effect is larger. The right-hand panel of Figure 6, however, illustrates that the backlash effect unambiguously magnifies the diversifying effect of the pro-diversity claim. The existence of a backlash effect thus suggests a trade-off whereby firms that employ pro-diversity claims may attract applicant pools that are smaller than those that employ non-discrimination claims, but are comprised of an even greater share of applicants from historically marginalized groups.

### **3.8 Additional Implications**

Besides these first-order implications for applicant pools size and composition, these findings also provide insight into analogous effects related to worker retention. It is well established that firms reduce turnover to the extent they align with the values of their workers (Chatman 1989, 1991, O'Reilly, Chatman and Caldwell 1991), with more recent research demonstrating the retention-enhancing role of fit on sociopolitical values (Bermiss and McDonald 2018). Given these findings, the same potential applicants who are deterred by feelings of group threat from applying to employers that make pro-diversity claims are unlikely to fit well in, and will be more inclined to leave, the increasingly racially diverse workplaces resulting from growing labor market diversity. Avoiding such workers may also reduce turnover among workers from marginalized groups by

reducing the risk of creating a work environment hostile to diversity. Since anti-diversity attitudes are difficult to observe via resumes or interviews but are predicted to become more common in response to growing diversity (Outten et al. 2012), this filtering feature may emerge as a strong motivation to use pro-diversity claims in recruiting efforts, even if they reduce applicant pool size.

This paper also advances literature regarding contingencies circumscribing the effects of non-pecuniary human capital strategies. To date, this research has mostly examined micro-level contingencies, including those at the organizational and individual-level. Examining individual-level contingencies, for example, Carnahan, Kryscynski and Olson (2017) illustrate how the relationship between corporate social responsibility and employee retention (Bode, Singh and Rogan 2015) is enhanced among employees who have experienced mortality-salient events, while Sauermann and Roach (2014) show that scientists' taste to pursue science (Stern 2004) is tempered by preferences for money. Examining organization-level contingencies, Abraham and Burbano (2021) demonstrate that the positive effect of pro-social activity on employee attraction (Turban and Greening 1997, Burbano 2016) varies with CEO gender. In examining labor market diversity, this paper, in contrast, joins recent work by Cotofan et al. (2021) in demonstrating that non-pecuniary workplace preferences, and by implication the effects of corresponding non-pecuniary human capital strategies, depend on the macro-level environment.

Finally, this paper advances broader literature regarding anti-diversity backlash. Past research shows how perceptions of growing diversity generate anti-diversity political preferences among members of dominant social groups (Craig and Richeson 2014*b*, Wetts and Willer 2018, Major, Blodorn and Blascovich 2018). These findings have garnered attention in academic (Mutz 2018, Jardina 2019) and popular (Resnick 2017, Edsall 2021) press for providing a partial explanation for the rise of national populist movements in the West. This paper provides the first evidence that

these processes extend beyond the political sphere, shaping not just how members of dominant groups vote, but also where they choose to work. Pro-diversity recruiting strategies may thus generate labor market sorting not just by race (as white workers increasingly eschew the pro-diversity employers that workers of color prefer), but also by racial attitudes (as those who feel most threatened by growing diversity increasingly spurn pro-diversity employers). Such patterns potentially reduce intergroup contact in a way that exacerbates exclusionary attitudes and political polarization.

### **3.9 Conclusion**

This paper challenges conventional wisdom that growing labor market diversity unambiguously enhances the applicant pool—expanding and diversifying advantages of pro-diversity over non-discrimination claims. Instead, I propose that growing labor market diversity may cause pro-diversity claims to have a contractionary effect on applicant pools by reducing attraction to these claims among applicants from dominant groups. I present pre-registered experimental evidence and exploratory archival analysis consistent with this “workplace backlash” effect. I develop a framework explaining conditions under which this effect leads pro-diversity claims to have a contractionary effect on applicant pools.

### 3.10 Works Cited

- Abascal, Maria. "Contraction as a Response to Group Threat: Demographic Decline and Whites' Classification of People Who Are Ambiguously White." *American Sociological Review* 85, no. 2 (April 1, 2020): 298–322. <https://doi.org/10.1177/0003122420905127>.
- Abraham, Mabel, and Vanessa Burbano. "Congruence Between Leadership Gender and Organizational Claims Affects the Gender Composition of the Applicant Pool: Field Experimental Evidence." *Organization Science* Articles in Advance (March 17, 2021): 1–25. <https://doi.org/10.1287/orsc.2021.1442>.
- Abrajano, Marisa, and Zoltan L. Hajnal. *White Backlash: Immigration, Race, and American Politics*. Princeton University Press, 2017.
- Apfelbaum, Evan P., Nicole M. Stephens, and Ray E. Reagans. "Beyond One-Size-Fits-All: Tailoring Diversity Approaches to the Representation of Social Groups." *Journal of Personality and Social Psychology* 111, no. 4 (2016): 547–66. <https://doi.org/10.1037/pspi0000071>.
- Barbulescu, Roxana, and Matthew Bidwell. "Do Women Choose Different Jobs from Men? Mechanisms of Application Segregation in the Market for Managerial Workers." *Organization Science* 24, no. 3 (June 1, 2013): 737–56. <https://doi.org/10.1287/orsc.1120.0757>.
- Bermiss, Y. Sekou, and Rory McDonald. "Ideological Misfit? Political Affiliation and Employee Departure in the Private-Equity Industry." *Academy of Management Journal* 61, no. 6 (December 1, 2018): 2182–2209. <https://doi.org/10.5465/amj.2016.0817>.
- Blalock, Hubert M. *Toward a Theory of Minority-Group Relations*. Wiley, 1967.
- Blumer, Herbert. "Race Prejudice as a Sense of Group Position." *The Pacific Sociological Review* 1, no. 1 (1958): 3–7. <https://doi.org/10.2307/1388607>.
- Bode, Christiane, Jasjit Singh, and Michelle Rogan. "Corporate Social Initiatives and Employee Retention." *Organization Science* 26, no. 6 (December 2015): 1702–20. <https://doi.org/10.1287/orsc.2015.1006>.
- Brundage, Vernon. "Profile Of The Labor Force By Educational Attainment." Spotlight on Statistics. U.S. Bureau of Labor Statistics, 2017.
- Burbano, Vanessa C. "Social Responsibility Messages and Worker Wage Requirements: Field Experimental Evidence from Online Labor Marketplaces." *Organization Science* 27, no. 4 (June 30, 2016): 1010–28. <https://doi.org/10.1287/orsc.2016.1066>.
- . "The Demotivating Effects of Communicating a Social-Political Stance: Field Experimental Evidence from an Online Labor Market Platform." *Management Science* 67, no. 2 (February 2021): 1004–25. <https://doi.org/10.1287/mnsc.2019.3562>.

- Campero, Santiago. "Hiring and Intra-Occupational Gender Segregation in Software Engineering." *American Sociological Review* 86, no. 1 (February 1, 2021): 60–92. <https://doi.org/10.1177/0003122420971805>.
- Carnahan, Seth, David Kryscynski, and Daniel Olson. "When Does Corporate Social Responsibility Reduce Employee Turnover? Evidence from Attorneys Before and After 9/11." *Academy of Management Journal* 60, no. 5 (October 2017): 1932–62. <https://doi.org/10.5465/amj.2015.0032>.
- Case, Anne, and Angus Deaton. "Rising Morbidity and Mortality in Midlife among White Non-Hispanic Americans in the 21st Century." *Proceedings of the National Academy of Sciences* 112, no. 49 (December 8, 2015): 15078. <https://doi.org/10.1073/pnas.1518393112>.
- Chatman, Jennifer A. "Improving Interactional Organizational Research: A Model of Person-Organization Fit." *Academy of Management Review* 14, no. 3 (July 1, 1989): 333–49. <https://doi.org/10.5465/amr.1989.4279063>.
- . "Matching People and Organizations: Selection and Socialization in Public Accounting Firms." *Administrative Science Quarterly* 36, no. 3 (1991): 459–84. <https://doi.org/10.2307/2393204>.
- Coff, Russell W. "Human Assets and Management Dilemmas: Coping with Hazards on the Road to Resource-Based Theory." *The Academy of Management Review* 22, no. 2 (1997): 374–402. <https://doi.org/10.2307/259327>.
- Cotofan, Maria, Lea Cassar, Robert Dur, and Stephan Meier. "Macroeconomic Conditions When Young Shape Job Preferences for Life." *The Review of Economics and Statistics*, May 6, 2021, 1–20. [https://doi.org/10.1162/rest\\_a\\_01057](https://doi.org/10.1162/rest_a_01057).
- Cox, Daniel, William A. Galston, and Robert P. Jones. "What It Means to Be an American: Attitudes in an Increasingly Diverse America Ten Years after 9/11." *Brookings* (blog), 2011. <https://www.brookings.edu/research/what-it-means-to-be-an-american-attitudes-in-an-increasingly-diverse-america-ten-years-after-911/>.
- Craig, Maureen A., and Jennifer A. Richeson. "More Diverse Yet Less Tolerant? How the Increasingly Diverse Racial Landscape Affects White Americans' Racial Attitudes." *Personality and Social Psychology Bulletin* 40, no. 6 (June 1, 2014): 750–61. <https://doi.org/10.1177/0146167214524993>.
- . "On the Precipice of a 'Majority-Minority' America: Perceived Status Threat From the Racial Demographic Shift Affects White Americans' Political Ideology." *Psychological Science* 25, no. 6 (June 1, 2014): 1189–97. <https://doi.org/10.1177/0956797614527113>.
- Cramer, K.J. *The Politics of Resentment: Rural Consciousness in Wisconsin and the Rise of Scott Walker*. Chicago Studies in American Politics. University of Chicago Press, 2016. <https://books.google.com/books?id=Rg2ZCwAAQBAJ>.

- Dobbin, F. *Inventing Equal Opportunity*. Princeton University Press, 2009.  
<https://books.google.com/books?id=GCAz25CZKoIC>.
- Dover, Tessa L., Brenda Major, and Cheryl R. Kaiser. “Members of High-Status Groups Are Threatened by pro-Diversity Organizational Messages.” *Journal of Experimental Social Psychology* 62 (January 1, 2016): 58–67. <https://doi.org/10.1016/j.jesp.2015.10.006>.
- Du Bois, W. E. B. *Black Reconstruction in America 1860-1880*. Simon and Schuster, 1935.
- Edsall, Thomas B. “Opinion | White Riot.” *The New York Times*, 2021, Jan. 13, 2021 edition, sec. Opinion. <https://www.nytimes.com/2021/01/13/opinion/capitol-riot-white-grievance.html>.
- “Educational Attainment in the United States: 2019.” United State Census Bureau, 2019. <https://www.census.gov/data/tables/2019/demo/educational-attainment/cps-detailed-tables.html>.
- “Educational Attainment Statistics.” Statistics Explained. Eurostat, 2020. [https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Educational\\_attainment\\_statistics](https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Educational_attainment_statistics).
- Fernandez, Roberto M., Emilio J. Castilla, and Paul Moore. “Social Capital at Work: Networks and Employment at a Phone Center.” *American Journal of Sociology* 105, no. 5 (2000): 1288–1356.
- Fernandez, Roberto M., and Isabel Fernandez-Mateo. “Networks, Race, and Hiring.” *American Sociological Review* 71, no. 1 (February 1, 2006): 42–71. <https://doi.org/10.1177/000312240607100103>.
- Flory, Jeffrey A., Andreas Leibbrandt, Christina Rott, and Olga Stoddard. “Increasing Workplace Diversity: Evidence from a Recruiting Experiment at a Fortune 500 Company.” *Journal of Human Resources* 57, no. 1 (2021): 73–92. <https://doi.org/10.3368/jhr.56.1.0518-9489R1>.
- Gest, Justin. *The New Minority: White Working Class Politics in an Age of Immigration and Inequality*. Oxford University Press, 2016.
- Gonyea, Don. “Majority Of White Americans Say They Believe Whites Face Discrimination.” You, Me And Them: Experiencing Discrimination In America. National Public Radio, October 24, 2017. <https://www.npr.org/2017/10/24/559604836/majority-of-white-americans-think-theyre-discriminated-against>.
- Gorbatai, Andreea, Peter Younkin, and Gordon Burtch. “Collateral Damage: The Relationship Between High-Salience Events and Variation in Racial Discrimination.” *Organization Science* Articles in Advance (December 15, 2021): 1–22. <https://doi.org/10.1287/orsc.2021.1527>.

- Highhouse, Scott, Filip Lievens, and Evan F. Sinar. “Measuring Attraction to Organizations.” *Educational and Psychological Measurement* 63, no. 6 (December 1, 2003): 986–1001. <https://doi.org/10.1177/0013164403258403>.
- Hochschild, Arlie Russell. *Strangers in Their Own Land: Anger and Mourning on the American Right*. The New Press, 2016.
- Jardina, A. *White Identity Politics*. Cambridge Studies in Public Opinion and Political Psychology. Cambridge University Press, 2019. <https://books.google.com/books?id=DUPXvAEACAAJ>.
- Jones, Robert P., Daniel Cox, E.J. Dionne, Jr., William A. Galston, Betsy Cooper, and Rachel Lienesch. “How Immigration and Concerns about Cultural Change Are Shaping the 2016 Election: Findings from the 2016 PRRI/Brookings Immigration Survey.” Public Religion Research Institute, 2016. <https://www.prrri.org/wp-content/uploads/2016/06/PRRI-Brookings-2016-Immigration-survey-report.pdf>.
- Kalev, Alexandra, Frank Dobbin, and Erin Kelly. “Best Practices or Best Guesses? Assessing the Efficacy of Corporate Affirmative Action and Diversity Policies.” *American Sociological Review* 71, no. 4 (August 2006): 589–617. <https://doi.org/10.1177/000312240607100404>.
- Kang, Sonia K., Katherine A. DeCelles, András Tilcsik, and Sora Jun. “Whitened Résumés: Race and Self-Presentation in the Labor Market.” *Administrative Science Quarterly* 61, no. 3 (September 1, 2016): 469–502. <https://doi.org/10.1177/0001839216639577>.
- Knowles, Eric D., Brian S. Lowery, Caitlin M. Hogan, and Rosalind M. Chow. “On the Malleability of Ideology: Motivated Construals of Color Blindness.” *Journal of Personality and Social Psychology* 96, no. 4 (April 2009): 857–69. <https://doi.org/10.1037/a0013595>.
- Leana, Carrie R., Vikas Mittal, and Emily Stiehl. “PERSPECTIVE—Organizational Behavior and the Working Poor.” *Organization Science* 23, no. 3 (June 1, 2012): 888–906. <https://doi.org/10.1287/orsc.1110.0672>.
- Leibbrandt, Andreas, and John A. List. “Do Equal Employment Opportunity Statements Backfire? Evidence From A Natural Field Experiment On Job-Entry Decisions.” National Bureau of Economic Research, September 17, 2018. <https://doi.org/10.3386/w25035>.
- Lin, Ken-Hou, Carolina Aragão, and J. Adam Cobb. “Women, Minorities, and Non-Union Workers Continue to Dominate Low-Wage Markets, and Experience Job Insecurity and Limited Upward Mobility.” *PRC Research Brief* 5, no. 10 (2020). <https://doi.org/10.26153/tsw/10795>.
- Maitner, Angela T., Eliot R. Smith, and Diane M. MacKie. “Intergroup Emotions Theory: Prejudice and Differentiated Emotional Reactions toward Outgroups.” In *The Cambridge Handbook of the Psychology of Prejudice*, edited by Chris G. Sibley and Fiona Kate Barlow, 111–30. Cambridge: Cambridge University Press, 2017. <https://doi.org/10.1017/9781316161579.006>.

- Major, Brenda, Alison Blodorn, and Gregory Blascovich. “The Threat of Increasing Diversity: Why Many White Americans Support Trump in the 2016 Presidential Election.” *Group Processes & Intergroup Relations* 21, no. 6 (September 1, 2018): 931–40. <https://doi.org/10.1177/1368430216677304>.
- Manstead, Antony S. R. “The Psychology of Social Class: How Socioeconomic Status Impacts Thought, Feelings, and Behaviour.” *The British Journal of Social Psychology* 57, no. 2 (April 2018): 267–91. <https://doi.org/10.1111/bjso.12251>.
- Margalit, Yotam. “Economic Insecurity and the Causes of Populism, Reconsidered.” *Journal of Economic Perspectives* 33, no. 4 (November 2019): 152–70. <https://doi.org/10.1257/jep.33.4.152>.
- McConnell, Christopher, Yotam Margalit, Neil Malhotra, and Matthew Levendusky. “The Economic Consequences of Partisanship in a Polarized Era.” *American Journal of Political Science* 62, no. 1 (January 1, 2018): 5–18. <https://doi.org/10.1111/ajps.12330>.
- Morath, Eric. “Worker Shortage Strains Return to Business.” *The Wall Street Journal*, 2021, May 7, 2021 edition. <https://www.wsj.com/articles/millions-are-unemployed-why-cant-companies-find-workers-11620302440>.
- Mutz, Diana C. “Status Threat, Not Economic Hardship, Explains the 2016 Presidential Vote.” *Proceedings of the National Academy of Sciences* 115, no. 19 (May 8, 2018): E4330–39. <https://doi.org/10.1073/pnas.1718155115>.
- Myers, Ella. “Beyond the Psychological Wage: Du Bois on White Dominion.” *Political Theory* 47, no. 1 (February 1, 2019): 6–31. <https://doi.org/10.1177/0090591718791744>.
- Norton, Michael I., and Samuel R. Sommers. “Whites See Racism as a Zero-Sum Game That They Are Now Losing.” *Perspectives on Psychological Science* 6, no. 3 (May 1, 2011): 215–18. <https://doi.org/10.1177/1745691611406922>.
- O’Reilly, Charles A., Jennifer Chatman, and David F. Caldwell. “People and Organizational Culture: A Profile Comparison Approach to Assessing Person-Organization Fit.” *Academy of Management Journal* 34, no. 3 (September 1, 1991): 487–516. <https://doi.org/10.5465/256404>.
- Outten, H. Robert, Michael T. Schmitt, Daniel A. Miller, and Amber L. Garcia. “Feeling Threatened About the Future: Whites’ Emotional Reactions to Anticipated Ethnic Demographic Changes.” *Personality and Social Psychology Bulletin* 38, no. 1 (January 2012): 14–25. <https://doi.org/10.1177/0146167211418531>.
- Pellegrin, Roland J., and Charles H. Coates. “Executives and Supervisors: Contrasting Definitions of Career Success.” *Administrative Science Quarterly* 1, no. 4 (1957): 506–17. <https://doi.org/10.2307/2390871>.
- Plaut, Victoria C., Flannery G. Garnett, Laura E. Buffardi, and Jeffrey Sanchez-Burks. “‘What about Me?’ Perceptions of Exclusion and Whites’ Reactions to Multiculturalism.”



- Journal of Personality and Social Psychology* 101, no. 2 (2011): 337–53.  
<https://doi.org/10.1037/a0022832>.
- Plaut, Victoria C., Kecia M. Thomas, Kyneshawau Hurd, and Celina A. Romano. “Do Color Blindness and Multiculturalism Remedy or Foster Discrimination and Racism?” *Current Directions in Psychological Science* 27, no. 3 (June 1, 2018): 200–206.  
<https://doi.org/10.1177/0963721418766068>.
- Purdie-Vaughns, Valerie, Claude M. Steele, Paul G. Davies, Ruth Dittmann, and Jennifer Randall Crosby. “Social Identity Contingencies: How Diversity Cues Signal Threat or Safety for African Americans in Mainstream Institutions.” *Journal of Personality and Social Psychology* 94, no. 4 (2008): 615–30. <https://doi.org/10.1037/0022-3514.94.4.615>.
- Quillian, Lincoln. “Prejudice as a Response to Perceived Group Threat: Population Composition and Anti-Immigrant and Racial Prejudice in Europe.” *American Sociological Review* 60, no. 4 (1995): 586–611. <https://doi.org/10.2307/2096296>.
- Ramarajan, Lakshmi, and Erin Reid. “Relational Reconciliation: Socializing Others across Demographic Differences.” *Academy of Management Journal* 63, no. 2 (April 1, 2020): 356–85. <https://doi.org/10.5465/amj.2017.0506>.
- Resnick, Brian. “White Fear of Demographic Change Is a Powerful Psychological Force.” *Vox*, 2017, Jan 28, 2017 edition. <https://www.vox.com/science-and-health/2017/1/26/14340542/white-fear-trump-psychology-minority-majority>.
- Richeson, Jennifer A, and Richard J Nussbaum. “The Impact of Multiculturalism versus Color-Blindness on Racial Bias.” *Journal of Experimental Social Psychology* 40, no. 3 (May 2004): 417–23. <https://doi.org/10.1016/j.jesp.2003.09.002>.
- Sauermann, Henry, and Michael Roach. “Not All Scientists Pay to Be Scientists: PhDs’ Preferences for Publishing in Industrial Employment.” *Research Policy* 43, no. 1 (February 1, 2014): 32–47. <https://doi.org/10.1016/j.respol.2013.07.006>.
- Snellman, Kaisa, and Peter Younkin. “Who’s the Boss? Evidence of Job Seeker Bias from a Field Experiment.” SSRN Scholarly Paper, February 14, 2021.  
<https://doi.org/10.2139/ssrn.3785311>.
- Stephens, Nicole M., Hazel Rose Markus, and L. Taylor Phillips. “Social Class Culture Cycles: How Three Gateway Contexts Shape Selves and Fuel Inequality.” *Annual Review of Psychology* 65, no. 1 (2014): 611–34. <https://doi.org/10.1146/annurev-psych-010213-115143>.
- Stern, Scott. “Do Scientists Pay to Be Scientists?” *Management Science* 50, no. 6 (June 1, 2004): 835–53. <https://doi.org/10.1287/mnsc.1040.0241>.
- Tajfel, H. “Social Psychology of Intergroup Relations.” *Annual Review of Psychology* 33, no. 1 (1982): 1–39. <https://doi.org/10.1146/annurev.ps.33.020182.000245>.

- Tavernise, Sabrina, and Robert Gebeloff. "Census Shows Sharply Growing Numbers of Hispanic, Asian and Multiracial Americans." *The New York Times*, August 12, 2021, sec. U.S.
- Taylor, Marylee C. "How White Attitudes Vary with the Racial Composition of Local Populations: Numbers Count." *American Sociological Review* 63, no. 4 (1998): 512–35. <https://doi.org/10.2307/2657265>.
- Toossi, Mitra. "A Century of Change: The U.S. Labor Force, 1950-2050." *Monthly Labor Review*, 2002, 14.
- Turban, Daniel B., and Daniel W. Greening. "Corporate Social Performance and Organizational Attractiveness to Prospective Employees." *The Academy of Management Journal* 40, no. 3 (1997): 658–72. <https://doi.org/10.2307/257057>.
- Wang, Heli C., and Jay B. Barney. "Employee Incentives to Make Firm-Specific Investments: Implications for Resource-Based Theories of Corporate Diversification." *The Academy of Management Review* 31, no. 2 (2006): 466–76. <https://doi.org/10.2307/20159212>.
- Wetts, Rachel, and Robb Willer. "Privilege on the Precipice: Perceived Racial Status Threats Lead White Americans to Oppose Welfare Programs." *Social Forces* 97, no. 2 (December 1, 2018): 793–822. <https://doi.org/10.1093/sf/soy046>.
- "Wide Gender Gap, Growing Educational Divide in Voters' Party Identification." Pew Research Center, 2018. <https://www.pewresearch.org/politics/2018/03/20/1-trends-in-party-affiliation-among-demographic-groups/>.
- Wilkins, Clara L., and Cheryl R. Kaiser. "Racial Progress as Threat to the Status Hierarchy: Implications for Perceptions of Anti-White Bias." *Psychological Science* 25, no. 2 (February 1, 2014): 439–46. <https://doi.org/10.1177/0956797613508412>.
- Williams, Margaret L., and Talya N. Bauer. "The Effect of a Managing Diversity Policy on Organizational Attractiveness." *Group & Organization Management* 19, no. 3 (September 1, 1994): 295–308. <https://doi.org/10.1177/1059601194193005>.
- Wilson, Valerie, Ethan Miller, and Melat Kassa. "Racial Representation in Professional Occupations: By the Numbers." Fact Sheet. Economic Policy Institute, 2021. <https://www.epi.org/publication/racial-representation-prof-occ/>.

### 3.11 Tables

Table 3.11.1 Subjects reported attraction to a firm that made a pro-diversity claim, a non-discrimination claim, or neither of these claims.

---

<i>Pro-Diversity:</i>	<b>Diversity:</b> We celebrate racial, cultural, and gender differences. [+ <i>Generic</i> ]
<i>Non-Discrimination:</i>	<b>Non-Discrimination:</b> We do not discriminate based on racial, cultural, or gender differences. [+ <i>Generic</i> ]
<i>Generic:</i>	Founded in 2016, we are known for working tirelessly to meet our clients' needs. We employ more than 150 employees across the country and are looking to expand.

---

Table 3.11.2 Among non-bachelor's-holding subjects, the white minority information significantly decreased attraction to the pro-diversity claim, generating a newfound preference for the non-discrimination claim.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	-0.18*** (0.04)	-0.12* (0.06)	-0.13* (0.06)	-0.09 (0.08)	-0.17* (0.08)	-0.17* (0.08)
Generic	-0.28*** (0.04)	-0.24*** (0.05)	-0.27*** (0.05)	-0.28*** (0.07)	-0.26*** (0.07)	-0.26*** (0.07)
No Bachelor's		-0.15** (0.05)				-0.21** (0.07)
Pro-Diversity × No Bachelor's		-0.10 (0.08)				0.09 (0.11)
Generic × No Bachelor's		-0.07 (0.07)				-0.02 (0.10)
White Min.			0.01 (0.05)	0.06 (0.07)	-0.05 (0.07)	-0.05 (0.07)
Pro-Diversity × White Min.			-0.10 (0.08)	-0.23* (0.10)	0.11 (0.11)	0.11 (0.11)
Generic × White Min.			-0.02 (0.07)	-0.05 (0.10)	0.04 (0.11)	0.04 (0.11)
White Min. × No Bachelor's						0.11 (0.10)
Pro-Diversity × White Min. × No Bachelor's						-0.34* (0.15)
Generic × White Min. × No Bachelor's						-0.10 (0.15)
Constant	2.87*** (0.03)	2.96*** (0.04)	2.86*** (0.04)	2.77*** (0.05)	2.98*** (0.05)	2.98*** (0.05)
Observations	2879	2879	2879	1696	1183	2879


Notes: The outcome in all models is organizational attraction. The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Robust standard errors in parentheses. Estimated with OLS. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

### 3.12 Figures

Figure 3.12.1 Subjects were randomly assigned to one of these six treatment conditions.

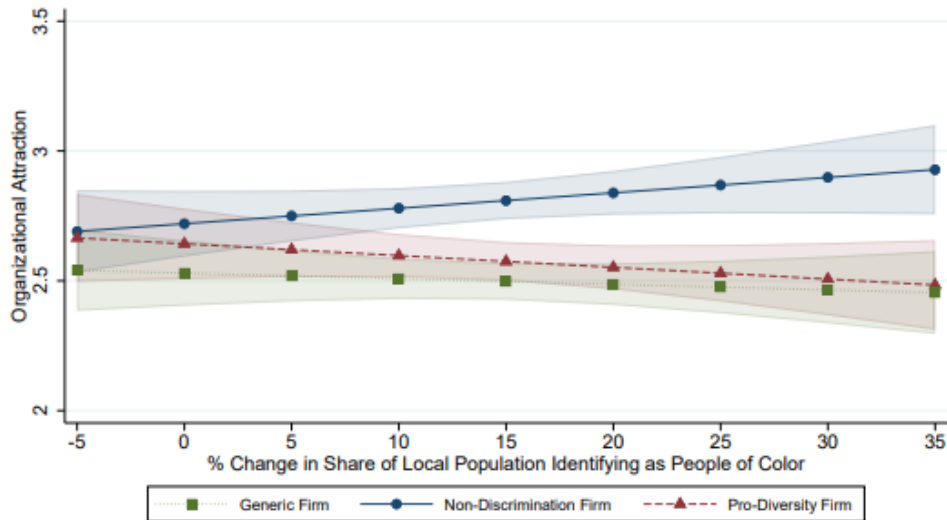
	Pro-Diversity Claim	Non-Discrimination Claim	Generic
White Majority Workforce	White Majority X Pro-Diversity Claim	White Majority X Non-Discrimination Claim	White Majority X Generic
White Minority Workforce	White Minority X Pro-Diversity Claim	White Minority X Non-Discrimination Claim	White Minority X Generic

Figure 3.12.2 Among non-bachelor's holding subjects, the white minority information significantly decreased attraction to the pro-diversity firm.



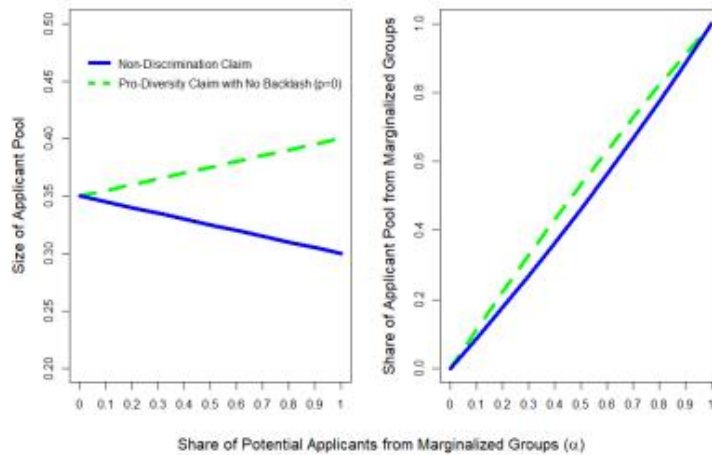
Notes: Bars represent mean organizational attraction by treatment condition by education status. Error bars indicate 95% confidence intervals. The p-value above each pair of bars is for the difference in the means represented by these bars. The left-hand panel excludes subjects that reported holding a bachelor's degree. The right-hand panel includes all other subjects. Refer to columns 4-6 of Table A14 for the regression results that correspond to this figure.

Figure 3.12.3 For non-bachelor's-holding subjects, larger increases in racial diversity in subjects' local labor markets were associated with a preference bifurcation in favor of non-discrimination claims.



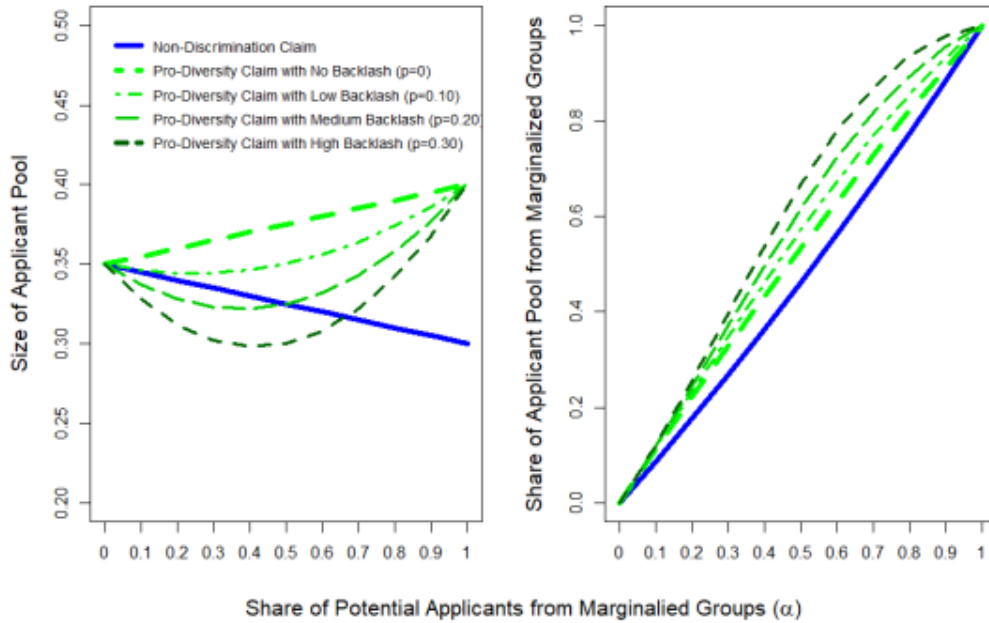
Notes: Points represent predicted values of ordinary least squares estimation of a version of Equation 2 where *WhiteMinority* is replaced with the percentage change between 2010 and 2019 in the share of residents in the subject's MSA of residence who identify as people of color. This is estimated on the sub-sample of non-bachelor's-holding subjects. Shaded regions represent 95% confidence intervals.

Figure 3.12.4 Without backlash, growing labor market diversity causes the pro-diversity claim to have an ever-larger expansionary effect on the applicant pool. The pro-diversity claim always results in greater applicant pool diversity.



Notes: The left panel illustrates applicant pool size for the two different claims as workers from marginalized groups come to comprise a larger share of the labor market. The right panel illustrates the share of the applicant pool from marginalized groups.

Figure 3.12.5 With backlash, growing labor market diversity often causes the pro-diversity claim to have a contractionary effect on the applicant pool. Backlash unambiguously enhances the diversifying advantages of the pro-diversity claim.



*Notes:* The left panel illustrates applicant pool size for the two different claims as workers from marginalized groups come to comprise a larger share of the labor market. The right panel illustrates the share of the applicant pool from marginalized groups.

### 3.13 Appendices

Table 3.13.1 Summary of subject characteristics by treatment condition shows that randomization was successful.

	White Majority			White Minority			
	Full Sample	Generic	Non-Discrimination Claim	Pro-Diversity Claim	Generic	Non-Discrimination Claim	Pro-Diversity Claim
Age	49.86 (16.1)	50.33 (16.11)	50.15 (16.11)	50.23 (16.11)	49.73 (16.11)	49.72 (16.11)	49.13 (16.11)
Female	0.51 (0.5)	0.53 (0.5)	0.52 (0.5)	0.52 (0.5)	0.49 (0.5)	0.5 (0.5)	0.5 (0.5)
High School or Less	0.22 (0.41)	0.19 (0.39)	0.21 (0.39)	0.23 (0.39)	0.2 (0.39)	0.24 (0.39)	0.25 (0.39)
Some College	0.24 (0.43)	0.24 (0.43)	0.23 (0.43)	0.25 (0.43)	0.28 (0.43)	0.19 (0.43)	0.26 (0.43)
Associates Degree	0.13 (0.34)	0.13 (0.34)	0.13 (0.34)	0.12 (0.34)	0.13 (0.34)	0.14 (0.34)	0.12 (0.34)
Bachelor's Degree	0.18 (0.39)	0.18 (0.38)	0.18 (0.38)	0.17 (0.38)	0.2 (0.38)	0.21 (0.38)	0.16 (0.38)
More than Bachelor's Degree	0.23 (0.42)	0.26 (0.44)	0.25 (0.44)	0.23 (0.44)	0.2 (0.44)	0.22 (0.44)	0.2 (0.44)
Democrat	0.31 (0.46)	0.32 (0.47)	0.33 (0.47)	0.33 (0.47)	0.31 (0.47)	0.28 (0.47)	0.29 (0.47)
Independent	0.34 (0.47)	0.34 (0.48)	0.34 (0.48)	0.34 (0.48)	0.34 (0.48)	0.37 (0.48)	0.32 (0.48)
Republican	0.35 (0.48)	0.34 (0.47)	0.33 (0.47)	0.33 (0.47)	0.36 (0.47)	0.35 (0.47)	0.39 (0.47)
Number of Observations	2879	471	455	448	480	509	516

Note: Standard deviation in parentheses. P-value of the null of no difference between the means of the indicated treatment condition and the White Majority X Generic Firm condition is reported in brackets.



Table 3.13.2 Comparison of mean organizational attraction by treatment condition.

	White Majority			
	Full Sample	Generic	Non-Discrimination Claim	Pro-Diversity Claim
Organizational Attraction	2.71 (0.84)	2.59 (0.8)	2.86 (0.78)	2.73 (0.89)
Number of Observations	2879	471	455	448
	White Minority			
	Full Sample	Generic	Non-Discrimination Claim	Pro-Diversity Claim
Organizational Attraction	2.71 (0.84)	2.58 (0.8)	2.88 (0.8)	2.64 (0.92)
Number of Observations	2879	480	509	516

*Note:* Standard deviation in parentheses.

Table 3.13.3 The results reported in Table A14 are unaffected by the inclusion of demographic controls.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	-0.176*** (0.038)	-0.116* (0.057)	-0.127* (0.055)	-0.099 (0.075)	-0.182* (0.080)	-0.173* (0.079)
Generic	-0.279*** (0.036)	-0.234*** (0.053)	-0.270*** (0.052)	-0.294*** (0.072)	-0.264*** (0.072)	-0.252*** (0.073)
No Bachelor's	-0.179*** (0.034)	-0.119* (0.057)	-0.177*** (0.034)			-0.169* (0.073)
Pro-Diversity × No Bachelor's		-0.102 (0.076)				0.078 (0.110)
Generic × No Bachelor's		-0.079 (0.072)				-0.032 (0.103)
White Min.			0.016 (0.051)	0.058 (0.069)	-0.048 (0.073)	-0.037 (0.073)
Pro-Diversity × White Min.			-0.094 (0.076)	-0.218* (0.102)	0.108 (0.113)	0.109 (0.113)
Generic × White Min.			-0.016 (0.072)	-0.043 (0.098)	0.039 (0.104)	0.035 (0.106)
White Min. × No Bachelor's						0.095 (0.101)
Pro-Diversity × White Min. × No Bachelor's						-0.336* (0.152)
Generic × White Min. × No Bachelor's						-0.090 (0.145)
Age	-0.004*** (0.001)	-0.004*** (0.001)	-0.004*** (0.001)	-0.000 (0.001)	-0.010*** (0.002)	-0.004*** (0.001)
Female	0.055 (0.033)	0.056 (0.033)	0.054 (0.033)	0.181*** (0.043)	-0.106* (0.049)	0.057 (0.033)
Republican	-0.175*** (0.032)	-0.175*** (0.032)	-0.173*** (0.032)	-0.180*** (0.041)	-0.159** (0.049)	-0.171*** (0.032)
Constant	3.202*** (0.055)	3.168*** (0.058)	3.195*** (0.061)	2.756*** (0.089)	3.541*** (0.089)	3.191*** (0.068)
Observations	2879	2879	2879	1696	1183	2879

Notes: The outcome in all models is organizational attraction. The excluded/comparison firm condition is the non-discrimination firm. The only difference between these estimates and those reported in Table A14 is the inclusion of controls for age, gender, and Republican political identity. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. All other estimates include all subjects. All models estimated with OLS. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.4 Subjects were more likely to agree or strongly agree that the pro-diversity firm, compared to the non-discrimination firm, was more likely than other companies to employ racial minorities. This perception did not change with exposure to information.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	0.058** (0.021)	0.019 (0.035)	0.056 (0.031)	0.112** (0.038)	-0.016 (0.050)	-0.016 (0.050)
Generic	-0.082*** (0.020)	-0.060 (0.033)	-0.065* (0.028)	-0.053 (0.033)	-0.085 (0.048)	-0.085 (0.048)
No Bachelor's		-0.126*** (0.030)				-0.198*** (0.043)
Pro-Diversity × No Bachelor's		0.073 (0.044)				0.127* (0.063)
Generic × No Bachelor's		-0.034 (0.041)				0.032 (0.058)
White Min.			0.033 (0.030)	0.090* (0.036)	-0.046 (0.048)	-0.046 (0.048)
Pro-Diversity × White Min.			0.004 (0.043)	-0.040 (0.053)	0.066 (0.070)	0.066 (0.070)
Generic × White Min.			-0.033 (0.040)	-0.080 (0.048)	0.046 (0.067)	0.046 (0.067)
White Min. × No Bachelor's						0.136* (0.060)
Pro-Diversity × White Min. × No Bachelor's						-0.106 (0.088)
Generic × White Min. × No Bachelor's						-0.126 (0.082)
Constant	0.301*** (0.015)	0.373*** (0.024)	0.284*** (0.021)	0.199*** (0.025)	0.397*** (0.035)	0.397*** (0.035)
Observations	2879	2879	2879	1696	1183	2879

Notes: The outcome in all models is a binary variable indicating whether or not the subject Agrees or Strongly Agrees that Company B is more likely than other companies to employ racial minorities. The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Estimated with linear probability models. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.5 Subjects were not more likely to agree or strongly agree that the pro-diversity firm, compared to the non-discrimination firm, was more likely more likely than other companies to have a boss who was a racial minority. Information did not change this.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	0.018 (0.021)	-0.006 (0.034)	-0.009 (0.029)	0.011 (0.035)	-0.029 (0.049)	-0.029 (0.049)
Generic	-0.076*** (0.019)	-0.081* (0.033)	-0.058* (0.028)	-0.064* (0.032)	-0.054 (0.047)	-0.054 (0.047)
No Bachelor's		-0.151*** (0.029)				-0.171*** (0.042)
Pro-Diversity × No Bachelor's		0.050 (0.042)				0.040 (0.061)
Generic × No Bachelor's		0.012 (0.040)				-0.010 (0.057)
White Min.			0.017 (0.029)	0.031 (0.035)	-0.006 (0.047)	-0.006 (0.047)
Pro-Diversity × White Min.			0.050 (0.041)	0.060 (0.050)	0.044 (0.069)	0.044 (0.069)
Generic × White Min.			-0.035 (0.039)	-0.010 (0.046)	-0.057 (0.065)	-0.057 (0.065)
White Min. × No Bachelor's						0.037 (0.059)
Pro-Diversity × White Min. × No Bachelor's						0.015 (0.085)
Generic × White Min. × No Bachelor's						0.047 (0.080)
Constant	0.277*** (0.014)	0.363*** (0.024)	0.268*** (0.021)	0.195*** (0.025)	0.366*** (0.035)	0.366*** (0.035)
Observations	2879	2879	2879	1696	1183	2879

Notes: The outcome in all models is a binary variable indicating whether or not the subject agrees or strongly agrees that Company B is more likely than other companies to have a racial minority boss. The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Estimated with linear probability models. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.6 Subjects were more likely to believe that the majority of employees within the pro-diversity firm, compared to the non-discrimination firm, were Democrats (rather than Republicans or Independents). This perception did not change with information.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	0.059** (0.022)	0.011 (0.034)	0.029 (0.032)	0.081 (0.043)	-0.042 (0.050)	-0.042 (0.050)
Generic	-0.165*** (0.023)	-0.123*** (0.035)	-0.180*** (0.032)	-0.217*** (0.042)	-0.132** (0.049)	-0.132** (0.049)
No Bachelor's		-0.054 (0.032)				-0.086 (0.046)
Pro-Diversity × No Bachelor's		0.082 (0.045)				0.123 (0.066)
Generic × No Bachelor's		-0.071 (0.045)				-0.085 (0.065)
White Min.			0.003 (0.032)	0.029 (0.042)	-0.032 (0.047)	-0.032 (0.047)
Pro-Diversity × White Min.			0.055 (0.044)	0.021 (0.057)	0.102 (0.069)	0.102 (0.069)
Generic × White Min.			0.029 (0.045)	0.045 (0.059)	0.015 (0.069)	0.015 (0.069)
White Min. × No Bachelor's						0.061 (0.064)
Pro-Diversity × White Min. × No Bachelor's						-0.081 (0.090)
Generic × White Min. × No Bachelor's						0.030 (0.091)
Constant	0.602*** (0.016)	0.632*** (0.024)	0.600*** (0.023)	0.563*** (0.031)	0.649*** (0.034)	0.649*** (0.034)
Observations	2879	2879	2879	1696	1183	2879

*Notes:* The outcome in all models is a binary variable indicating whether or not subjects believe the boss of Company B is a Democrat (rather than a Republican or an Independent). The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Estimated with linear probability models. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.7 Subjects were more likely to believe that the boss of the pro-diversity firm, compared to the boss of the non-discrimination firm, was a Democrat (rather than a Republican or Independent). This perception did not change with information.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	0.079*** (0.022)	0.044 (0.034)	0.042 (0.033)	0.078 (0.043)	-0.004 (0.050)	-0.004 (0.050)
Generic	-0.207*** (0.022)	-0.175*** (0.035)	-0.213*** (0.032)	-0.250*** (0.041)	-0.166*** (0.049)	-0.166*** (0.049)
No Bachelor's		-0.095** (0.032)				-0.109* (0.046)
Pro-Diversity × No Bachelor's		0.064 (0.045)				0.082 (0.066)
Generic × No Bachelor's		-0.053 (0.045)				-0.085 (0.064)
White Min.			-0.019 (0.032)	-0.009 (0.043)	-0.035 (0.048)	-0.035 (0.048)
Pro-Diversity × White Min.			0.068 (0.045)	0.055 (0.059)	0.093 (0.069)	0.093 (0.069)
Generic × White Min.			0.011 (0.045)	0.044 (0.058)	-0.022 (0.069)	-0.022 (0.069)
White Min. × No Bachelor's						0.026 (0.064)
Pro-Diversity × White Min. × No Bachelor's						-0.038 (0.090)
Generic × White Min. × No Bachelor's						0.066 (0.090)
Constant	0.561*** (0.016)	0.615*** (0.024)	0.571*** (0.023)	0.525*** (0.031)	0.634*** (0.035)	0.634*** (0.035)
Observations	2879	2879	2879	1696	1183	2879

Notes: The outcome in all models is a binary variable indicating whether or not subjects believe the boss of Company B is a Democrat (rather than a Republican or an Independent). The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Estimated with linear probability models. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.8 Subjects did not perceive the pro-diversity firm as more or less likely than the non-discrimination firm to unfairly discriminate against white workers. Information did not change this.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	-0.028 (0.022)	-0.012 (0.034)	-0.052 (0.032)	-0.069 (0.041)	-0.027 (0.050)	-0.027 (0.050)
Generic	-0.077*** (0.021)	-0.008 (0.034)	-0.096** (0.031)	-0.165*** (0.039)	-0.006 (0.049)	-0.006 (0.049)
No Bachelor's		-0.004 (0.031)				-0.003 (0.046)
Pro-Diversity × No Bachelor's		-0.025 (0.044)				-0.041 (0.065)
Generic × No Bachelor's		-0.119** (0.043)				-0.159* (0.063)
White Min.			-0.055 (0.031)	-0.056 (0.041)	-0.053 (0.047)	-0.053 (0.047)
Pro-Diversity × White Min.			0.046 (0.043)	0.059 (0.056)	0.027 (0.069)	0.027 (0.068)
Generic × White Min.			0.035 (0.042)	0.074 (0.054)	-0.010 (0.067)	-0.010 (0.067)
White Min. × No Bachelor's						-0.002 (0.063)
Pro-Diversity × White Min. × No Bachelor's						0.032 (0.089)
Generic × White Min. × No Bachelor's						0.083 (0.086)
Constant	0.356*** (0.015)	0.358*** (0.024)	0.385*** (0.023)	0.383*** (0.030)	0.387*** (0.035)	0.387*** (0.035)
Observations	2879	2879	2879	1696	1183	2879

Notes: The outcome in all models is a binary variable indicating whether or not the subject agrees or strongly agrees that Company B is more likely than other companies to unfairly discriminate against white workers. The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Estimated with linear probability models. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.9 Subjects did not perceive the pro-diversity firm as more or less likely than the non-discrimination firm to unfairly favor racial minority workers. Information did not change this.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	-0.022 (0.022)	0.035 (0.034)	-0.037 (0.032)	-0.098* (0.041)	0.054 (0.050)	0.054 (0.050)
Generic	-0.068** (0.021)	-0.024 (0.033)	-0.081** (0.031)	-0.105* (0.041)	-0.049 (0.047)	-0.049 (0.047)
No Bachelor's		0.004 (0.031)				0.018 (0.046)
Pro-Diversity × No Bachelor's		-0.092* (0.044)				-0.152* (0.065)
Generic × No Bachelor's		-0.076 (0.043)				-0.056 (0.062)
White Min.			-0.043 (0.031)	-0.055 (0.041)	-0.027 (0.047)	-0.027 (0.047)
Pro-Diversity × White Min.			0.028 (0.043)	0.077 (0.055)	-0.038 (0.069)	-0.038 (0.069)
Generic × White Min.			0.023 (0.042)	0.010 (0.055)	0.049 (0.066)	0.049 (0.066)
White Min. × No Bachelor's						-0.028 (0.062)
Pro-Diversity × White Min. × No Bachelor's						0.115 (0.089)
Generic × White Min. × No Bachelor's						-0.039 (0.086)
Constant	0.349*** (0.015)	0.346*** (0.023)	0.371*** (0.023)	0.379*** (0.030)	0.361*** (0.035)	0.361*** (0.035)
Observations	2879	2879	2879	1696	1183	2879

Notes: The outcome in all models is a binary variable indicating whether or not the subject agrees or strongly agrees that Company B is more likely than other companies to unfairly favor racial minority workers. The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Estimated with linear probability models. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .



Table 3.13.10 Subjects did not perceive the pro-diversity firm as a fairer place to work than the non-discrimination firm. Information did not change this.

	(1) All	(2) All	(3) All	(4) No Bachelor's	(5) Bachelor's+	(6) All
Pro-Diversity	-0.016 (0.022)	-0.030 (0.035)	-0.003 (0.033)	0.045 (0.042)	-0.063 (0.052)	-0.063 (0.052)
Generic	-0.129*** (0.021)	-0.094** (0.034)	-0.160*** (0.031)	-0.179*** (0.037)	-0.139** (0.049)	-0.139** (0.049)
No Bachelor's		-0.124*** (0.032)				-0.164*** (0.046)
Pro-Diversity × No Bachelor's		0.033 (0.045)				0.108 (0.066)
Generic × No Bachelor's		-0.057 (0.043)				-0.040 (0.062)
White Min.			-0.051 (0.031)	-0.020 (0.040)	-0.095 (0.049)	-0.095 (0.049)
Pro-Diversity × White Min.			-0.024 (0.044)	-0.075 (0.056)	0.061 (0.071)	0.061 (0.071)
Generic × White Min.			0.059 (0.042)	0.053 (0.052)	0.084 (0.069)	0.084 (0.069)
White Min. × No Bachelor's						0.075 (0.063)
Pro-Diversity × White Min. × No Bachelor's						-0.135 (0.090)
Generic × White Min. × No Bachelor's						-0.030 (0.086)
Constant	0.384*** (0.016)	0.454*** (0.024)	0.411*** (0.023)	0.341*** (0.029)	0.505*** (0.036)	0.505*** (0.036)
Observations	2879	2879	2879	1696	1183	2879

Notes: The outcome in all models is a binary variable indicating whether or not the subject agrees or strongly agrees that Company B is a fair place to work. The excluded/comparison firm condition is the non-discrimination firm. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 include only bachelor's-holding subjects. Estimated with linear probability models. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.11 Non-bachelor's-holding subjects residing in labor markets that experienced greater percentage increases in racial diversity from 2010 to 2-10 were significantly less attracted to the pro-diversity firm.

	(1) All	(2) No Bachelor's	(3) Bachelor's+	(4) All
Pro-Diversity	-0.103 (0.070)	-0.085 (0.089)	-0.144 (0.119)	-0.144 (0.119)
Generic	-0.225*** (0.067)	-0.209* (0.088)	-0.258** (0.100)	-0.258** (0.100)
Diversity Change	-0.001 (0.003)	0.005 (0.004)	-0.010* (0.005)	-0.010* (0.005)
Pro-Diversity × Diversity Change	-0.006 (0.004)	-0.010* (0.005)	0.003 (0.008)	0.003 (0.008)
Generic × Diversity Change	-0.004 (0.004)	-0.007 (0.005)	0.001 (0.007)	0.001 (0.007)
No Bachelor's				-0.351*** (0.097)
Pro-Diversity × No Bachelor's				0.059 (0.148)
Generic × No Bachelor's				0.049 (0.133)
No Bachelor's × Diversity Change				0.015* (0.006)
Pro-Diversity × No Bachelor's × Diversity Change				-0.013 (0.010)
Generic × No Bachelor's × Diversity Change				-0.008 (0.008)
Constant	2.888*** (0.048)	2.741*** (0.062)	3.091*** (0.075)	3.091*** (0.075)
Observations	2647	1541	1106	2647

Notes: The outcome in all models is organizational attraction. "Diversity Change" indicates percentage change between 2010 and 2019 in share of local population in the subject's MSA of residence that identify as people of color. Column 2 includes only non-bachelor's holding subjects. Column 3 includes only bachelor's holding subjects. Estimated with OLS. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ .

Table 3.13.12 Additional test show the robustness of the finding that non-bachelor's-holding subjects residing in labor markets that experienced greater percentage increases in the share of residents of color between 2010 & 2019 were less attracted to pro-diversity.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Pro-Diversity	-0.079 (0.089)	-0.090 (0.091)	-0.077 (0.090)	-0.092 (0.091)	-0.087 (0.091)	-0.079 (0.089)	-0.101 (0.091)	-0.076 (0.090)	-0.082 (0.091)	-0.073 (0.091)
Generic	-0.212* (0.088)	-0.201* (0.087)	-0.196* (0.088)	-0.215* (0.088)	-0.220* (0.088)	-0.217* (0.092)	-0.216* (0.093)	-0.200* (0.094)	-0.212* (0.093)	-0.214* (0.093)
Diversity Change	0.005 (0.004)	0.006 (0.004)	0.006 (0.004)	0.004 (0.004)	0.004 (0.004)	0.007 (0.004)	0.005 (0.005)	0.006 (0.005)	0.006 (0.004)	0.007 (0.004)
Pro-Diversity × Diversity Change	-0.011* (0.005)	-0.009 (0.005)	-0.010* (0.005)	-0.009 (0.005)	-0.010 (0.005)	-0.011* (0.005)	-0.009 (0.005)	-0.011* (0.005)	-0.011* (0.005)	-0.012* (0.005)
Generic × Diversity Change	-0.007 (0.005)	-0.007 (0.005)	-0.008 (0.005)	-0.006 (0.005)	-0.007 (0.005)	-0.007 (0.005)	-0.006 (0.005)	-0.008 (0.005)	-0.006 (0.005)	-0.007 (0.005)
Constant	2.715*** (0.099)	2.858*** (0.352)	2.843*** (0.355)	130.484** (43.678)	132.916** (44.559)	2.668*** (0.108)	3.065*** (0.505)	2.981*** (0.515)	121.013* (48.744)	118.801* (49.389)
Individual Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
MSA Controls		✓	✓	✓	✓	✓	✓	✓	✓	✓
% Change MSA Controls				✓	✓	✓	✓	✓	✓	✓
State Fixed Effects						✓	✓	✓	✓	✓
Observations	1541	1541	1541	1541	1541	1541	1541	1541	1541	1541

Notes: These are robustness specifications for the results reported in column 2 of Table A11. All estimates include only non-bachelor's-holding subjects. The outcome in all models is organizational attraction. "Diversity Change" indicates percentage change between 2010 and 2019 in share of local population in the subject's MSA that identify as people of color. The excluded/comparison firm condition is the non-discrimination firm. The coefficient of interest is "Pro-Diversity x Diversity Change." "Individual Controls" include the subject's age, gender, and whether or not they identify as a Republican. "MSA Controls" include the share of residents in the MSA for both 2010 and 2019 that hold a bachelor's degree, the MSA's population, the median age within the MSA, and average earnings within the MSA. "% Change MSA Controls" is simply the percent change for each of these variables between 2010 and 2019. All models estimated with OLS. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . In the four cases with no stars on the coefficient of interest (columns 2, 4, 5, and 7) the p-value is 0.06, or just barely over the conventional 0.05 threshold.

Table 3.13.13 Analysis of performance on an optional editing task.

	(1) All	(2) All	(3) All	(4) Non-Bachelor's	(5) Bachelor's +	(6) All
Pro-Diversity	0.10 (0.11)	-0.08 (0.16)	-0.02 (0.16)	0.09 (0.21)	-0.17 (0.26)	-0.17 (0.26)
Generic	0.17 (0.11)	-0.07 (0.16)	-0.02 (0.16)	0.20 (0.21)	-0.31 (0.23)	-0.31 (0.23)
No Bachelor's		-0.01 (0.15)				0.01 (0.23)
Pro-Diversity × No Bachelor's		0.29 (0.21)				0.26 (0.33)
Generic × No Bachelor's		0.40 (0.22)				0.51 (0.32)
White Min.			-0.29* (0.14)	-0.30 (0.19)	-0.27 (0.22)	-0.27 (0.22)
Pro-Diversity × White Min.			0.23 (0.21)	0.25 (0.28)	0.17 (0.33)	0.17 (0.33)
Generic × White Min.			0.36 (0.22)	0.26 (0.29)	0.48 (0.33)	0.48 (0.33)
White Min. × No Bachelor's						-0.03 (0.29)
Pro-Diversity × White Min. × No Bachelor's						0.08 (0.43)
Generic × White Min. × No Bachelor's						-0.22 (0.43)
Constant	1.00*** (0.07)	1.00*** (0.11)	1.15*** (0.11)	1.15*** (0.15)	1.14*** (0.18)	1.14*** (0.18)
Observations	2879	2879	2879	1696	1183	2879

*Notes:* The outcome in all models is the number of corrections made in an optional editing task. The excluded/comparison firm condition is the non-discrimination firm. The only difference between these estimates and those reported in Table A14 is that subgroup analysis is conducted by reported partisanship. Estimates in Column 4 only include non-bachelor's-holding subjects. Estimates in column 5 includes bachelor's-holding subjects. All other estimates include all subjects. All models estimated with OLS. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . This analysis was pre-registered, and is reported here for completeness. See the addendum to the pre-analysis plan (PreAnalysisPlanDeviations.pdf) for an explanation as to why this analysis is not included in the body of the text.

Table 3.13.14 Analysis of treatment effects by political partisanship.

	(1) All	(2) All	(3) All	(4) Republican	(5) Non-Republican	(6) All
Pro-Diversity	-0.18*** (0.04)	-0.10* (0.05)	-0.13* (0.06)	-0.20* (0.09)	-0.09 (0.07)	-0.09 (0.07)
Generic	-0.28*** (0.04)	-0.32*** (0.05)	-0.27*** (0.05)	-0.22** (0.08)	-0.30*** (0.07)	-0.30*** (0.07)
Republican		-0.18*** (0.05)				-0.15 (0.08)
Pro-Diversity × Republican		-0.19* (0.08)				-0.11 (0.11)
Generic × Republican		0.09 (0.07)				0.08 (0.11)
White Min.			0.01 (0.05)	-0.02 (0.08)	0.04 (0.06)	0.04 (0.06)
Pro-Diversity × White Min.			-0.10 (0.08)	-0.16 (0.12)	-0.04 (0.10)	-0.04 (0.10)
Generic × White Min.			-0.02 (0.07)	-0.01 (0.11)	-0.03 (0.09)	-0.03 (0.09)
White Min. × Republican						-0.06 (0.10)
Pro-Diversity × White Min. × Republican						-0.13 (0.16)
Generic × White Min. × Republican						0.02 (0.15)
Constant	2.87*** (0.03)	2.94*** (0.03)	2.86*** (0.04)	2.78*** (0.06)	2.92*** (0.05)	2.92*** (0.05)
Observations	2879	2879	2879	1696	1183	2879

*Notes:* The outcome in all models is organizational attraction. The excluded/comparison firm condition is the non-discrimination firm. The only difference between these estimates and those reported in Table A14 is that subgroup analysis is conducted by reported partisanship. Estimates in Column 4 only include subjects who identified as Republican. Estimates in column 5 includes those that identified as Democrats or Independents. All other estimates include all subjects. All models estimated with OLS. Robust standard errors in parentheses. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . This analysis was pre-registered, and is reported here for completeness. See the addendum to the pre-analysis plan (PreAnalysisPlanDeviations.pdf) for an explanation as to why this analysis is not included in the body of the text.

### 3.13.1 Qualtrics Methodology for Removing Low-Quality Respondents

As mentioned in the text, Qualtrics removed low-quality respondents. This was done through their proprietary, standardized “data scrub” procedure. This was completed solely by Qualtrics. My only participation was to request the scrub, which I did before the completion of data collection, before analyzing data, and after Qualtrics indicated that there appeared to be a non-trivial number of “low quality” respondents. Upon my request, Qualtrics provided the following list of techniques used to identify low-quality respondents. This is the exact description provided by Qualtrics:

1. Survey responses: Our scrubbing process identifies bad respondents based on information provided in the survey. Respondent’ survey responses are used to evaluate if the person was paying attention while taking the survey or was there for some malicious intent

2. Patterns: patterns help in identifying the underlying behavior while completing the survey. Patterns also help in evaluating whether the survey responders are reading the questions or not. E.g. Flatliners

3. Click thru behavior: The way responders click through the survey is a great way for identifying presence of mind. Responders who simply click through the questions without reading are ones who generally display a flatliner behavior or are ones who take the survey in abnormally fast speed. All behaviors indicator of lack of attention while taking the survey

4. Duplicate responses: Identification of duplicate responses without using PII is a key feature of the scrubbing offering. Based on responders’ demographics, attitudinal responses, keyboard typing habits, and verbatim responses, we can identify duplicates within the survey. People can take surveys from multiple machines, multiple emails, multiple locations etc. however, evaluating their style of responding and survey data ensures their responses are not included in analysis.

5. Keystroke analysis: Human beings are habitual creatures, analyzing how they type using keyboard is a great way of identifying duplicate response behavior.

6. Machine Responses (Bots): Using a mix of duplicates, keystroke analysis, verbatim responses, our scrubbing algorithms can identify the presence of any machine-based responses

7. Inattentiveness: Identifying respondents who take too long to complete the survey or respond in a contradictory manner is critical in ensuring results are not biased.

## **Chapter 4 The Effect of Flatter Hierarchy on Applicant Pool Gender Diversity: Evidence from Experiments**

with Saerom Lee and Justin Frake

### **Abstract**

This paper investigates how job seekers' perceptions of an employer's formal hierarchy affect the size and gender composition of its applicant pool. Building on the literature on gendered organizations and organizational design, we develop opposing perspectives on these relationships. To arbitrate between these perspectives, we first conduct a field experiment in partnership with a hiring firm. We find that featuring a flatter hierarchy in recruiting materials does not significantly affect the size of the applicant pool, but significantly decreases women's representation within it. Our follow-up survey experiment identifies several potential mechanisms (e.g., perceptions of career progression, informality, workload, and fit). Our findings imply that firms' growing tendency to adopt flatter hierarchies could inadvertently undermine efforts to attract a greater proportion of women applicants.



## 4.1 Introduction

Women remain chronically underrepresented in many labor market contexts (Fernandez and Campero 2017, Murciano-Goroff 2021, Phillips 2005). As the likelihood of an employer hiring a woman is typically proportional to women's representation in its applicant pool (Fernandez and Abraham 2011, Fernandez-Mateo and Fernandez 2016, Petersen et al. 2000), the literature on gender-based labor market segregation has considered various recruitment strategies whereby employers might ameliorate this under-representation by increasing the share of women applicants. In particular, prior research illustrates that employers can augment women's representation by highlighting specific organizational characteristics that women job seekers, compared to men job seekers, find disproportionately attractive, such as evidence of promoting women and commitments to diversity or social causes (Abraham and Burbano 2022, Flory et al. 2019, Wiswall and Zafar 2018).

In this paper, we examine how women's representation in an employer's applicant pool may be shaped by its formal hierarchy, which is a fundamental feature of all formal organizations that is represented by the number of management levels (Burton and Obel 2004, Puranam 2018, Simon 1997). Although the extensive literature on organizational design has investigated how a firm's hierarchical structure shapes the day-to-day experiences and long-term career progression of employees *within* the firm (e.g., Acker 1990, Baron 1984, Lee and Edmondson 2017), little research has examined how its hierarchical structure may be perceived by job seekers and thus affect their propensity to apply to the firm in the first place—let alone, whether this propensity varies by gender. Understanding the relationship between an employer's hierarchical structure and its applicant pool is theoretically as well as phenomenologically important, given that employers commonly reference their hierarchical structures in their recruiting efforts, as we document in Figure 1.

To examine this relationship, we draw upon prior literature on gendered organizations and

organizational design to develop opposing perspectives regarding how job seekers perceive an employer's organizational structure and how these perceptions shape their inclination to apply. We first discuss the perspective that job seekers may be more attracted to a *flatter* hierarchy with fewer management levels, because they may perceive this organizational structure to satisfy a need for autonomy and flexibility and to act as “an egalitarian alternative” to traditional, taller hierarchies (Foss and Klein 2014:73, Oldham and Hackman 1981, Reitzig 2022). This attraction, moreover, may be disproportionately stronger among women job seekers inasmuch as prior research argues they may perceive traditional, taller hierarchies as perpetuating men's advantage (Acker 1990, Ferguson 1984) and serving as “an important location of male dominance” (Acker 1990:139) with “masculine principles dominating their authority structures” (Kanter 1977:46). If women tend to hold this negative perception of taller hierarchies and/or if they disproportionately value autonomy, flexibility, and egalitarian work environments (Barbulescu and Bidwell 2013, Bartling et al. 2014, Deci and Ryan 2000), they may be relatively more attracted than men to flatter employers.

We then elaborate the opposing perspective that job seekers may be more attracted to a *taller* hierarchy with more management levels. For one, job seekers may perceive that, because these organizational structures clarify responsibilities and facilitate accountability (Lee 2022, Puranam 2018, Tetlock 1985), taller hierarchies are more likely than flatter hierarchies to satisfy their unconscious need for order and stability (Friesen et al. 2014, Fromm 1941, Tiedens et al. 2007) and equitable rewards distribution (Gruenfeld and Tiedens 2010, Halevy et al. 2011). Moreover, they may perceive taller organizational structures as more likely to fulfill their desire for career progression by offering more opportunities for promotion (Baron et al. 1986, Murray 1938). In contrast, job seekers may be repelled not only by the perception that flatter employers provide fewer promotional opportunities, but also that this scarcity will result in more intense competition

and conflict (Baron et al. 1986, Flory et al. 2015). This perception may be stronger among women job seekers, who are, compared to men, less likely to prefer or succeed in workplaces with intense competition and conflict (Barrymore et al. 2022, Flory et al. 2015, Niederle and Vesterlund 2011). More generally, women may be disproportionately disinclined to apply to flatter employers if they perceive these structures as having less managerial oversight and cultivating informal “bro cultures” that marginalize women (Chang 2018, Kanter 1977).

We evaluate these opposing perspectives by conducting a pair of experiments in the context of the U.S. labor market. First, partnering with a U.S. healthcare startup, we ran a pre-registered large-scale field experiment, in which we manipulated more than 8,200 job seekers’ perceptions of the hiring firm’s hierarchical structure. We find that characterizing the firm’s organizational structure as “flatter” in its recruiting efforts has, on average, no statistically significant effect on job seekers’ propensity to show interest in or apply to the open position. However, we find that, compared to men, women job seekers are significantly less likely to show interest in or apply to a flatter employer. More specifically, characterizing our partner firm’s organizational structure as “flatter” reduces the proportion of interested women job seekers by 14% and lowers the share of women applicants by 28%. To replicate this finding and explore the underlying mechanisms, we conducted a pre-registered large-scale survey experiment of 9,000 subjects where we measure how an employer’s hierarchy shapes job seekers’ attraction and their perceptions of a wide range of organizational characteristics. We again find that, compared to men, women are less attracted to flatter hierarchies. In addition, we find that this relative aversion may reflect the fact that, compared to men, women perceive flatter structures to (1) offer less opportunity for career advancement, (2) burden them with more work, and (3) be more difficult to fit into.

Our study makes several contributions. First, we advance the literature on labor market sorting

and segregation by gender by addressing recent calls to examine how cross-gender differences in perceptions of organizational traits may lead women and men to sort into different employers (Fernandez-Mateo and Kaplan 2018:1229). We do so by investigating how women and men job seekers differ in their perceptions of and preferences for formal hierarchy and, consequently, how they vary in their propensity to supply their labor to employers with different hierarchical structures. Though our results confirm Acker's (1990) insight that "organizational structure is not gender-neutral" (p.139), we challenge the idea that enhancing women's representation will be achieved by creating "non-hierarchical, egalitarian organizations" (p.141).

Second, we contribute to the literature on organizational design by unveiling how perceptions of a firm's formal hierarchy affects its human capital acquisition. To date, this literature has mainly focused on how hierarchical structure motivates and coordinates existing employees (for reviews, see Burton and Obel 2004, Oldham and Hackman 1981, Puranam 2018), neglecting its effects on attracting new employees. We demonstrate that flatter organizations may not only fail to attract more applicants but also decrease gender diversity in the applicant pool.

Third, by examining our question in the context of a high-growth startup, we speak to the ongoing debate regarding the appropriate organizational structure for new ventures. In contrast to received wisdom in entrepreneurship research that startups should have flatter hierarchies (Burns and Stalker 1961), we find startups with flatter hierarchies may struggle to attract a gender-diverse applicant pool. Lastly, we answer recent calls to elucidate the selection processes underlying entrepreneurial team formation (Shah et al. 2019) by assessing how individuals self-select into an entrepreneurial venture based on their perceptions of its formal hierarchy.

## 4.2 Background literature

### 4.2.1 Labor market sorting by gender

Given the persistent under-representation of women in many labor market contexts (Bielby and Baron 1986, Fernandez and Campero 2017, Murciano-Goroff 2021, Petersen and Morgan 1995, Reskin 1993, Reskin and Padavic 1994), a rich stream of research has explored why women and men may systematically sort into different employers (Blau et al. 1998, Blau and Hendricks 1979, Phillips 2005). These studies have primarily taken one of two perspectives to explain gender segregation (Fernandez and Sosa 2005, Fernandez-Mateo and Kaplan 2018). One perspective focuses on the characteristics of job seekers (i.e., the supply-side of the labor market): that is, how men and women job seekers could have different skills and preferences which cause them to select into different employers (Barbulescu and Bidwell 2013, Wiswall and Zafar 2018). The other perspective, in contrast, centers on the attributes of employers (i.e., the demand-side of the labor market): how employers' gender-biased or discriminatory characteristics (in particular, processes of employee selection, promotion, and retention) could prevent women from entering or remaining in certain firms (Bielby and Baron 1986, Perry et al. 1994, Reskin and Roos 1990).

This study takes up recent calls to move beyond this dichotomy (Fernandez-Mateo and Kaplan 2018) by investigating how women and men job seekers differ in how they *perceive* employers' demand-side characteristics (i.e., organizational traits) and how these differences in *perceptions* affect their selection into applicant pools. As Barbulescu and Bidwell (2013:741) point out, this approach to understanding gender-based labor market segregation “lies at the intersection of supply and demand: although segregation is engendered by workers' application decisions, those decisions anticipate the expected behavior of employers.” This integrated approach better accounts for the reality that women workers are often in very high

demand, especially among firms and industries in which they are least represented (Lambrecht and Tucker 2019).

From this integrative perspective, under-representation of women often arises because women job seekers perceive a greater likelihood of career-limiting gender bias based on certain organizational characteristics, such as an employer's leaders (e.g., their gender, political ideology, organizational blueprint, Campero and Kacperczyk 2020, Carnahan and Greenwood 2018, Phillips 2005), recruiting activities (e.g., the use of gendered language in its job postings, Castilla and Rho 2023, Gaucher et al. 2011), work arrangements (e.g., work-life balance, team-based work structures, Barbulescu and Bidwell 2013, Kalev 2009, Mas and Pallais 2017, Wiswall and Zafar 2018), reward systems (e.g., meritocracy, pay formalization, Abraham 2017, Castilla 2008), and even third-party resource providers (Abraham 2020). As job seekers have limited information about potential employers, they perceive these organizational characteristics through various means, such as job postings, online reviews, and conversations with current employees (DeVaro 2005, Dineen and Allen 2016). Job ads play a particularly important role in shaping workers' perceptions of the firm (Barber and Roehling 1993) because they not only provide detailed information about the employer and the open position but also commonly serve as the online location through which job seekers submit their applications. Accordingly, firms can curate job postings with the goal that they will be perceived more favorably by women and thus attract an applicant pool with a greater share of women (Abraham and Burbano 2022, Castilla and Rho 2023, Mihaljević et al. 2022).

#### 4.2.2 Formal hierarchy

Taking the integrative perspective on workplace gender segregation, we examine how men and women job seekers perceive an employer's formal hierarchy and how these perceptions may

affect their decision to apply. The literature on organizational design defines formal hierarchy as the vertical division of tasks represented by the layers of management in an organization (Burton and Obel 2004:75–77, Puranam 2018:106–126, Simon 1997:7). By adding managerial levels and dividing decision-making responsibilities across these levels, taller hierarchical structures reduce the number of subordinates that each manager directly supervises (Graicunas 1937, Csaszar 2021, Lee 2022). This smaller span of control allows boundedly rational managers to pay more attention to each of their subordinates, better integrate and evaluate the efforts of their subordinates, and resolve conflicts among them (Colombo and Grilli 2013, Puranam 2018:113). By fostering coordination and accountability, these multi-layered structures can establish order and stability within the organization (Simon 1997, Tetlock 1985).

In contrast, by removing managerial layers, flatter hierarchies reduce status differentiation, increase the span of control of managers and distribute more decision-making responsibilities to their subordinates (Graicunas 1937, Csaszar 2021, Lee 2022). As subordinates are each subject to less managerial supervision and are more empowered in flatter organizations, they can enjoy more autonomy and flexibility in allocating their efforts toward tasks that they perceive to match their skills and/or to be more meaningful (Oldham and Hackman 1981, Reitzig 2022, Saxenian 1996:76). This increased autonomy and flexibility, along with less status differentiation, can help create a sense of fairness and “egalitarian” culture (Foss and Klein 2014:73), where employees recognize each other as equals (Siegel et al. 2013:1174) and more freely voice their opinions (Keum and See 2017).

As formal hierarchy is such a fundamental feature of all organizations (Blau and Scott 2003, Burton and Obel 2004, Puranam 2018, Simon 1997), employers commonly highlight their hierarchical structures in their job postings. For example, panel (a) of Figure 1 displays a recent

job posting in which the employer underscores its “very flat hierarchy” with two managerial levels. We document the prevalence of such job postings by assessing the word occurrence of terms related to formal hierarchy (e.g., “hierarchical structure”, for more detail, see the footnote in panel (b)) in roughly 157 million job postings from more than 6 million employers in the U.S. between 2010 and 2019. Panel (b) shows that about 10% of our sample (i.e., 0.6 million employers) mentioned their hierarchical structures in their job postings. Though this proportion has remained relatively consistent over the last decade, the share of employers specifically highlighting a flatter organizational structure has roughly doubled. In line with this growing trend, flatter hierarchies have increasingly gained attention from academics and practitioners (e.g., Burton et al. 2017, Foss 2003, Foss and Klein 2022, Lee and Edmondson 2017, Lee 2022, Puranam and Hakonsson 2015, Reingold 2016, Reitzig 2022, Sorenson 2022).

Despite this heightened scholarly interest, little research has investigated how formal hierarchy affects applicant attraction. To address this gap, we develop opposing perspectives regarding how organizational structure may shape the size and gender composition of the applicant pool. In the following two sections, we first elaborate how job seekers, on average, perceive a firm’s hierarchical structure and then outline how these perceptions may vary by gender.

#### 4.2.3 How perceptions of hierarchy may affect the size of the applicant pool

On the one hand, job seekers may generally prefer flatter employers. By imposing fewer layers of management, flatter organizations decrease managerial oversight and status differentiation, while increasing autonomy and flexibility (Oldham and Hackman 1981, Lee 2022, Saxenian 1996:76). Although individual preferences vary, a growing body of work suggests that workers, on average, desire and value autonomy and flexibility in the workplace (Bartling et al. 2014, Deci and Ryan 2000, Spector 1986). These attributes cause workers to experiment and direct their efforts towards



tasks that they perceive to be more meaningful (Keum and See 2017, Lee and Edmondson 2017, Lee 2022). In addition, job seekers may perceive flatter employers as more fair and egalitarian because they reduce status differentiation among employees, distribute power more evenly, and ameliorate in-group biases (Foss and Klein 2014:73, Puranam 2018, Reitzig and Sorenson 2013). These perceptions potentially cause flatter organizations to attract more applicants.

On the other hand, job seekers may prefer taller employers. By adding managerial levels, taller hierarchies reduce the number of subordinates that each manager directly supervises (Csaszar 2021, Graicunas 1937), thereby clarifying responsibilities and facilitating accountability (Puranam 2018:113, Tetlock 1985). Job seekers may perceive that these clearly defined roles, responsibilities, and accountability satisfy their unconscious need for order and stability (Friesen et al. 2014, Fromm 1941, Tiedens et al. 2007). Furthermore, because more hierarchical levels represent more opportunities for promotion and career progression (Baron et al. 1986), job seekers may anticipate that taller organizations are more likely to fulfill their desire for career advancement (Murray 1938). In addition, because fewer subordinates are vying for the attention of each manager who makes promotion decisions (Halac and Prat 2016), job seekers may anticipate taller employers as being less competitive, a feature that candidates tend to prefer (Flory et al. 2015). Lastly, because the smaller span of control makes it easier for managers to evaluate and distribute rewards based on employees' efforts and contributions, rather than rationing those rewards equally, job seekers may perceive taller structures to be more equitable (Gruenfeld and Tiedens 2010, Halevy et al. 2011). Given these perceptions, taller organizations may attract more applicants.

#### 4.2.4 How perceptions of hierarchy may affect the gender composition of the applicant pool

In the previous section, we discussed how job seekers may, on average, be more or less attracted to flatter hierarchies. However, prior work has shown that women and men often differ in

their perceptions of various organizational characteristics (Barbulescu and Bidwell 2013, Croson and Gneezy 2009, Wiswall and Zafar 2018). If women and men job seekers also differ in their perceptions of an employer's formal hierarchy, these differences likely shape the gender composition of the employer's applicant pool. Hence, we theorize how perceptions of formal hierarchy may vary by gender and how this variation may lead to systemic differences in the gender composition of the employer's applicant pool.

Compared to men, women may be more attracted to flatter employers than taller ones because they perceive flatter hierarchies as ameliorating the gender biases they perceive as inherent to taller hierarchies. Feminist scholars have characterized taller hierarchies as representing a "scientific organization of inequality" (Ferguson 1984:7) which creates "a veneer of fairness" (Nicholson 2010), conceals "a gendered substructure" (Acker 1990:154), and reinforces patriarchal power structures (Reskin 1988). From this perspective, taller organizational structures impose narrow criteria for career progression which do not account for the societal expectations that women should tend to domestic responsibilities (Glass and Estes 1997, Nelson and Bridges 1999, Tomaskovic-Devey 1993). Thus, relative to men, women may be concerned that in taller organizations, they will be forced to choose between fulfilling domestic responsibilities or pursuing professional goals (Glass and Estes 1997). In contrast, women may expect that, by removing managerial layers that calcify gendered assumptions, flatter employers will provide more autonomy and flexibility that allow them to balance these responsibilities. Accordingly, prior studies have shown that women are more likely to apply to jobs offering flexibility regarding when and where the job is performed (He et al. 2021) and that, compared to men, they are willing to give up seven times more salary for such flexibility (Wiswall and Zafar 2018). In addition to flexibility and autonomy, women may more generally perceive flatter organizations as egalitarian work environments where women are more likely to be

treated fairly and succeed. In this regard, Ridgeway (2011:175) argues that flatter hierarchies benefit women “who seek equal outcomes with their male colleagues” because they offer more equal access to resources and opportunities while allowing women to avoid “bad actors” with biased perceptions and agendas. Supporting this argument, prior studies find that women are more likely to prefer equality and egalitarianism, whereas men prefer when rewards are allocated according to individual contributions (Brockner and Adsit 1986, Kahn et al. 1980).

Contrary to this perspective, however, women may be *less* inclined than men to apply to flatter employers due to the perception that less clear roles, responsibilities, and accountability in these organizations will enable the emergence of a male-dominated work environment. In this vein, Kanter (1977) and Chang (2018) characterize the work environment in flatter organizations as “old boys’ clubs” or “bro cultures.” Furthermore, large-scale surveys (e.g., Vassallo et al. 2017) and numerous firsthand accounts from women who have worked for flatter employers (e.g., Diamond 2019, Finley 2014, Flower Horne 2014, Hamburger 2021, Hunt 2017, Mont 2017) illustrate how men in this environment can saddle women with more work, especially lower value, less-interesting tasks that are unlikely to lead to promotion (e.g., cleaning, note-taking, food delivery). In addition to the work environment, the perception of greater competition for fewer promotions may be disproportionately unattractive to women, who often fare worse in and seek to avoid competition (Barrymore et al. 2022, Flory et al. 2015, Niederle and Vesterlund 2011) partly because, compared to men, they tend to undervalue their work and avoid self-promotion (Babcock and Laschever 2009, Exley and Kessler 2019, Lerchenmueller et al. 2019). In contrast, by delineating more hierarchical levels, reducing each manager’s span of control, and clarifying roles and responsibilities, taller organizations can reduce competition over promotion and managerial recognition. Thus, women job seekers may perceive taller hierarchies as a “great leveler” that safeguards against marginalization (Baron et al. 2007).

#### 4.2.5 Theoretical summary

In outlining these contrasting perspectives, we highlight important theoretical and practical tension. On the one hand, job seekers may prefer to work for *flatter* employers because they perceive that these firms offer more autonomy, flexibility, and egalitarian work environments than their taller counterparts. Because women tend to value these characteristics more than men, flatter organizations may be disproportionately attractive to women job seekers. On the other hand, job seekers may find *taller* employers more attractive because they perceive these organizations to provide more clarity about roles and responsibilities, more equitable treatment, and more opportunities for career advancement. If women expect these attributes of taller hierarchies to reduce competition and shield them from an informal “bro culture,” they may be especially attracted to taller hierarchies. In what follows, we empirically adjudicate between these opposing perspectives.

### **4.3 Methodological approach**

To identify the effect of an employer’s formal hierarchy on its applicant pool, we implement two complementary labor market experiments, both of which were approved by the Institutional Review Board and pre-registered.<sup>2</sup> First, partnering with a U.S. healthcare startup, we run a field experiment to identify our effects of interest in a context with high external validity. Second, as it is difficult to measure mechanisms in the field experiment (Chatterji et al. 2016), we conduct a follow-up survey experiment to explore mechanisms as well as replicate our field-study findings.

### **4.4 Study 1: Field experiment**

#### 4.4.1 Sampling

For our field experiment, we partnered with a U.S. healthcare startup, which allowed us to recruit for two positions: a software engineer and a business development representative. For these positions, we identified a pool of job candidates on ZipRecruiter, a widely-used online job search platform. Using this platform’s search function, we found candidates who had updated their profiles within

the last 30 days and completed at least a bachelor's degree. For the software engineer position, we limited our search to candidates who included "software engineer" in their past or current job titles and indicated "Healthcare" or "Engineering" as their areas of interest. For the business development position, we limited the search to those who included "business development" in their past or current job titles and indicated "Healthcare" or "Sales and Biz Dev" as their areas of interest. After applying these criteria, we gathered the first 4,200 individuals returned in each search (sorted by the date they last updated their profile) and thus identified a total of 8,400 job seekers.

#### 4.4.2 Experimental manipulation

As in previous reverse audit studies (e.g., Abraham and Burbano 2022, Flory et al. 2019), we embedded our experimental manipulation in the emails that we sent once to each of the 8,400 potential applicants. These emails consisted of four paragraphs (see Figure A1 in Appendix A1). The email's second paragraph featured the randomized manipulation, which was included for the treatment group (i.e., *Flatter* condition) but excluded for the control group (i.e., *No Information* condition). Adapting the job posting in panel (a) of Figure 1, this manipulation stated that the company has "a flat organizational structure" with "fewer levels of management than similarly sized startups in our industry." This statement was intended to reduce the number of hierarchical levels that subjects perceived the firm to have, without introducing other confounding characteristics.

The control condition for the *Flatter* treatment could have been operationalized in several other ways. One approach would have been to describe a *taller* hierarchy with "more management levels than similarly sized startups in our industry." But, this approach is problematic because firms rarely state such a description in their job postings (i.e., 0.0003% of companies in our analysis for Figure 1) and thus subjects may have negatively responded to the abnormality of the statement.

Another approach would have been to use an "inert" organizational description of a similar length

that did not reference formal hierarchy. However, any description of the organization may have not been inert, and instead may have had an unintended direct effect on the participant's attraction to the employer. For these reasons, we elected to use a control condition that made no mention of the hierarchical structure. We address the concerns that our treatment effects may have been driven by the difference in lengths of the treatment or the mere mention of hierarchy in Study 2 (see Section 5.7), where we demonstrate that the *Flatter* condition is, in fact, perceived by subjects to have fewer levels of management than the *No Information* condition. The follow-up experiment also includes a *Taller* condition to rule out the explanation that our results are driven by the mere mention of the organization's hierarchy.

The remaining paragraphs were identical for both treatment and control groups. In the first paragraph, we invited job seekers to apply for one of the positions. We designed the third paragraph to make sure job seekers in both conditions had the same perceptions of the firm's founding year, number of employees, and intention to grow. We took this measure in light of prior research suggesting that organizations with flatter hierarchies may be perceived as exceptionally young (Stinchcombe 1965), small (Burton and Obel 2004:168), or have no intention to grow (Lee and Kim 2022). The last paragraph included a hyperlink to the partner company's application web page, which subjects were invited to visit if they were interested in applying. This web page was identical for all subjects, regardless of their assigned treatment conditions.

The success of this field experiment depended on job seekers receiving and opening the emails we sent. To maximize delivery rates, we used the email delivery service Mailgun.com, which more than 225 thousand businesses employ for their mass-email marketing. Using this service, we sent an email once to each subject, tracked whether the email was successfully delivered, and removed from the sample 136 subjects that did not receive the email. Among the remaining 8,264 email

recipients, some subjects could have decided not to open our emails under any circumstance (i.e., “never-takers”, Angrist and Pischke 2008:158). However, because subjects decide whether to open the email before observing the randomized treatment (which is embedded in the body of the email and thus unobservable before opening the email), non-compliance should be unrelated to treatment status.<sup>5</sup> Technically, since we cannot remove never-takers, we estimate an “intention-to-treat” effect (ITT, i.e., the effect of being sent the email characterizing the employers as flatter), rather than the “average treatment effect” (ATE, i.e., the effect of actually perceiving the hiring employer as flatter). Because the never-takers will never show interest in or apply to the job, the ITT is smaller in magnitude than the ATE. Hence, our field experiment represents a conservative approach to estimating our theorized relationship.

#### 4.4.3 Measurement

**Dependent variables:** We measured job seekers’ attraction to the partner firm in two ways. First, *Click<sub>i</sub>* captures whether subject *i* clicked the hyperlink, embedded at the bottom of the email, to the partner company’s application web page.<sup>6</sup> We used this measure for two main reasons. First, because the treatment and the hyperlink are both embedded in the email, visiting the application website is the first behavioral outcome that we can observe in terms of the subject’s attraction.<sup>7</sup> Second, this outcome is unlikely to be affected by other information about the company, which subjects may observe after visiting the company’s application web page but before deciding whether to apply. Despite these two benefits, this measure may not accurately reflect a subject’s interest in applying to the company. For example, even if they are not interested, subjects in the *Flatter* condition might click on the hyperlink because they are simply curious to learn about a flatter organization. To address this concern, we also measure *Apply<sub>i</sub>*, which indicates whether subject *i* submitted an application via our partner company’s web page.

**Independent variable:** Our treatment variable  $Flatter_i$  captures whether subject  $i$  received an email stating that the organization’s hierarchy is flatter than its competitors’ (see Figure A1).

**Individual attributes:** As ZipRecruiter does not collect information on job seekers’ gender (presumably to reduce gender discrimination by employers), we used the website Gender API’s machine-learning algorithm, which used subjects’ full name to predict their gender ( $Womani$ ). In addition, to compare the effect of  $Flatter$  within (but not between) jobs, we added a binary variable ( $Software_i$ ) for whether the subject was invited to the software engineer position as opposed to the business development position. To assess whether our randomization procedures achieved a balance between the treatment and control groups, we also measured job seekers’ characteristics using self-reported information on ZipRecruiter. These attributes include the subjects’ years of work experience, most recent job title, most recent employer, highest level of education (i.e., undergraduate, master’s, or doctorate), and undergraduate institution.

#### 4.4.4 Summary statistics and randomization balance

Among the 8,400 emails we sent to the job seekers, 136 were not delivered because these emails were blocked by the subjects’ email servers. In addition, Gender API’s algorithm was unable to predict the gender of 97 job seekers.<sup>10</sup> Thus, our final sample was 8,167 job seekers. The first column of Table 1 shows the summary statistics across all subjects.<sup>11</sup> Of these 8,167 subjects, 28.5% showed interest in the open positions by clicking on the hyperlink, while 9.7% applied. Approximately 26.9% of the subjects were women, and 31.8% held a managerial position in their most recent job. In terms of the highest level of education, 35.8% had a master’s degree, while 1.7% had completed a doctorate. In turn, the second and third columns provide the summary statistics by treatment condition. The differences in observable characteristics of candidates between conditions are small and insignificant, suggesting that randomization was successful and that our estimates can be interpreted as causal.



Finally, the fourth and fifth columns show the summary statistics by position.

#### 4.4.5 Estimation

To assess the effect of a flatter hierarchy on attraction, we estimate the pre-registered equation:

$$Attraction_i = \beta_0 + \beta_1 Flatter_i + \beta_2 Software_j + \varepsilon_i \quad (1)$$

where  $Attraction_i$  stands for one of the two dependent variables (i.e.,  $Click_i$  or  $Apply_i$ ).  $Software_j$  is a dummy variable equal to one if the advertised position is for the “software engineer” position and zero for the “business development” position. The term  $\beta_1$  represents the average treatment effect of the *Flatter* condition relative to the *No Information* condition. The terms  $\beta_0$  and  $\varepsilon_i$  are the intercept and the random error term, respectively. In turn, the heterogeneous treatment effects with respect to gender are estimated using the following pre-registered equation:

$$Attraction_i = \beta_0 + \beta_1 Flatter_i + \beta_2 Woman_i + \beta_3 Flatter_i \times Woman_i + \beta_4 Software_j + \varepsilon_i \quad (2)$$

where  $\beta_3$  is the difference in the effect of the *Flatter* condition for women relative to men. We estimate Equations 1 and 2 Using OLS, which is equivalent to ANOVA when using a binary independent variable.<sup>13</sup> Recent research highlights that fixed-effects estimators utilize between-group variation when interaction effects are present (Shaver 2019, Giesselmann and Schmidt-Catran 2022). In response, we also employ double-demeaned fixed effects estimators as suggested by Giesselmann and Schmidt-Catran (2022), along with split sample analysis recommended by Shaver (2019). We also use robust standard errors to address the possibility that errors are correlated within a job.

#### 4.4.6 Results

The results are reported in Table 2 and displayed using bar graphs in Figure 2. In Table 2, panel (a) presents the mean differences using the full sample. In this panel, Models 1 to 5 measure attraction in terms of whether subjects showed interest in the open positions by clicking on the email’s hyperlink

to the company's application web page (i.e., *Click*), while Models 6 to 10 measure whether subjects applied to the job (i.e., *Apply*). Panels (b) and (c) repeat the analysis from panel (a) for only those subjects who were recruited for each job opening (business development and software).

We begin by examining the effect of a flatter hierarchy on subjects' propensity to click. In panel (a), Model 1 shows that the estimated treatment effect is small, negative, and indistinguishable from zero ( $p = .764$ ). This null result could imply one of the following. First, it may suggest that our treatment failed to manipulate the participants' perceptions of hierarchy. However, we believe this is unlikely because Study 2, which uses a similar treatment, shows that the *Flatter* condition is perceived by subjects to have fewer levels of management than the *No Information* condition (see Section 5.7). Second, given that our manipulation very likely succeeded, this finding may indicate that a flatter hierarchy does not impact in any meaningful way the subjects' propensity to click on the hyperlink. Lastly, this insignificant result may mean that a flatter hierarchy does affect the subjects' propensity to click but elicits offsetting responses. For example, as argued in Section 2.4, a flatter hierarchy may increase men's propensity to click but decrease women's propensity to click. We explore this explanation and find supportive evidence below.

In line with the notion that employers with flatter hierarchies attract disproportionately fewer women job seekers, Model 2 indicates that, compared to men, women subjects were five percentage points less likely to click the hyperlink in the *Flatter* condition compared to the *No Information* condition ( $p = .024$ ). In light of the fact that fixed-effects estimators rely on between-group variation when interaction effects are present (Shaver 2019, Giesselmann and Schmidt-Catran 2022), we conduct a parallel analysis using double-demeaned estimators, as recommended by Giesselmann and Schmidt-Catran (2022), in Model 5. The results are very similar, suggesting that between-job variation is not biasing the results in Model 2. To assess this divergence in the propensity to click,

Models 4 and 5 estimate the treatment effect of a flatter hierarchy by splitting the observations by gender. Model 4 reveals that women were approximately four percentage points less likely to click in the *Flatter* condition compared to the *No Information* condition ( $p = .036$ ). In contrast, Model 5 shows that men were one percentage point more likely to click in the *Flatter* condition but the 95% confidence interval overlaps zero. Overall, we document, that while 28% of those who clicked the hyperlink in the *No Information* condition (i.e., 324 out of 1,167) were women, only 24% of those who clicked in the *Flatter* condition (i.e., 277 out of 1,160) were women—hence, a reduction of 14%.

In Models 6 to 10, we examine the effect of a flatter hierarchy on subjects' propensity to apply for the job. In line with Model 1, Model 5 exhibits no main effect of a flatter hierarchy. However, like Models 2 to 5, Models 7 to 10 support the prediction that a flatter hierarchy decreases women's representation in the applicant pool. Models 7 and 8 reveal that, compared to men, women were about 3.5 percentage points less likely to apply in the *Flatter* condition than in the *No Information* condition ( $p = .017$  and  $.011$ , respectively). Similarly, Model 9 shows that women were three percentage points less likely to apply in the *Flatter* condition than in the *No Information* condition ( $p = .018$ ). In contrast, Model 10 indicates that men were about half a percentage point more likely to apply in the *Flatter* condition but the 95% confidence interval overlaps zero. Whereas 27% of the applicants (i.e., 108 out of 403) were women in the *No Information* condition, only 19% (i.e., 76 out of 392) were women in the *Flatter* condition—thus, a reduction of 28%.

Taken together, these results suggest that a flatter organizational structure does not substantively affect the size of the applicant pool, but significantly decreases women's representation within it.

#### 4.4.7 Post-hoc analyses of heterogeneity across job positions

Before progressing to Study 2, we offer exploratory, non-pre-registered analyses that explore potential heterogeneous effects across the two job positions. These analyses not only shed light on underlying

mechanisms but also help address concerns about fixed-effects estimators utilizing between-group variation when interaction terms are present. First, these two occupations differ in many job-level characteristics. For example, compared to the business development (BD) occupation, the software engineering (SW) occupation tends to be more technical and knowledge-intensive, have more modularized tasks, offer more autonomy, and require less coordination and in-person interactions, while having lower representation of women and more insular culture.<sup>14</sup> These differences could drive women's and men's relative attraction to flatter hierarchies in opposing ways. On the one hand, this relative attraction could be higher for the SW position than for the BD position because, compared to the latter, the former could be perceived as providing more autonomy and requiring less coordination and in-person interactions. On the other hand, it could be higher because the SW job might be perceived as having fewer women and a more insular culture than the BD job.

In addition to their job-level characteristics, these occupations may consist of different types of job seekers, as job seekers typically select their occupations based on their unique skills and preferences. For example, men and women seeking a SW position are likely to have a specific educational background (notably, a bachelor's degree in computer science or engineering), a specialized skill set (e.g., computer programming), and a stronger preference for autonomy, flexibility, and other non-pecuniary rewards (Stern 2004). In contrast, those seeking a BD position may have a broader range of educational backgrounds (including a bachelor's degree in business, economics, psychology, and other social sciences), a more general skill set (e.g., communication and soft skills), and a greater preference for interpersonal relationships, promotion, and pecuniary incentives (Bennett 2013). Thus, these differences in individual-level characteristics could also lead to across-position heterogeneity.<sup>15</sup>

Because these positions differ in many, non-randomized ways both at the occupation level and

at the job-seeker level and because we did not randomize subjects across positions, it is not possible to credibly infer why the effect might differ from one job to the other. Nonetheless, examining across-position heterogeneity might provide suggestive evidence of mechanisms possibly driving these effects. Hence, in Table 2, we replicate the analysis in panel (a) but limit the sample to the subjects recruited for the BD position (panel (b)) and the SW position (panel (c)), respectively (for bar charts, see panels (e) to (h) of Figure 2). Across both subsamples, we again find no main effect of the *Flatter* condition on the outcome *Click* or *Apply*, but the point estimates of the interaction between *Flatter* and *Woman* are negative for both outcomes. Interestingly, for the outcome *Click*, the point estimate of the interaction effect is larger in magnitude for the BD position ( $\hat{\beta} = -0.07$ ,  $p = .01$ ) than for the SW position ( $\hat{\beta} = -0.03$ ,  $p = .4$ ). In contrast, for *Apply*, the point estimate is smaller in magnitude for the BD position ( $\hat{\beta} = -0.03$ ,  $p = .06$ ) than for the SW position ( $\hat{\beta} = -0.05$ ,  $p = .07$ ). For neither outcome, however, is the difference in the effect sizes across the two jobs statistically different from zero.<sup>16</sup> Hence, these split sample analyses demonstrate that the main results in panel (a) are driven by similar responses by the applicant pools for both positions.

#### **4.5 Study 2: Survey experiment**

We complement our field study with a pre-registered survey experiment on Amazon’s Mechanical Turk (MTurk). This follow-up experiment had three objectives. First, it allowed us to examine whether our main finding regarding heterogeneous preference by gender would replicate in a broader population of subjects (e.g., those without college degrees or with lower-paying jobs). Second, it enabled us to address the concern that Study 1’s results are driven not by perceptions of how flat the organization’s hierarchy is, but by novelty aversion, the difference in the email length, or the mere mention of a flatter hierarchy. We did so by (1) including a *Taller* condition in addition to the *Flatter* and *No Information* conditions and (2) performing a manipulation check to measure the

number of hierarchical levels that the subjects perceived in these conditions. Lastly, the survey experiment allowed us to explore the mechanisms that we were unable to examine in the field study.

One common concern with using MTurk is whether its workers (MTurkers) are representative of U.S. workers. Recent studies (e.g., Difallah et al. 2018, Moss et al. 2020, Snowberg and Yariv 2021) have shown that MTurkers are similar to the U.S. population in terms of occupational and racial composition, but that they tend to be younger, more likely to be women, more educated, and more likely to be employed. We also observe these comparable demographic attributes in our 8,498 subjects (see Appendix A6). Although their median household income tends to be less than that of the U.S. population (e.g., \$47,000 vs. \$57,000, Difallah et al. 2018), most characterize their MTurk work as paid leisure (56%) and a financial source for non-essential expenses (69%). Only 8% consider it a full-time job (Moss et al. 2020). Despite these differences, Snowberg and Yariv (2021) demonstrate that MTurk workers and a representative sample of the U.S. population show similar comparative statics and correlations between behaviors. Furthermore, Kees et al. (2017) find that, compared to other survey data sources (e.g., student and professional panel samples), MTurk offers similar or more reliable data.

#### 4.5.1 Sampling

In Study 2, we posed as a human resources analytics company and recruited approximately 9,000 MTurk workers in the U.S.<sup>17</sup> These subjects were asked to complete a task entitled “Give feedback on recruiting material,” which described a job posting for a part-time, remote copy-editor position. We chose a remote position for two reasons: (1) to increase the realism of the experiment, as the subjects are already working remotely and our intended sample size was too large to be geographically focused, and (2) to increase the generalizability of our findings to non-traditional workers. Remote work, however, may be a conservative setting for examining our research question because the

perceived effects of hierarchy may be muted when subjects do not expect in-person interactions with coworkers or may not aspire to be promoted.

#### 4.5.2 Experimental manipulations

As in Study 1, we embedded experimental manipulations in the recruiting material, which consists of four sections as shown in Figure A2 in Appendix A7. The first section described the open position (i.e., “part-time, remote copyeditor”). The second section (“About Us”) provided information on the company. The third and fourth sections listed the job responsibilities and requirements, respectively. Here, the first, third, and fourth sections remained identical across all conditions.

The second section included our manipulations, which we present in detail in Figure A3 in Appendix A8. As in our field experiment, the *Flatter* condition stated that the firm has a “flat organizational structure” with “fewer levels of management than similarly sized startups in our industry.” For the *No Information* condition, this statement was omitted. As discussed in Section 4.2, we chose to provide no information as a control condition in Study 1 because any attempt to include an “inert” condition that described some organizational characteristic might itself have an unintended direct effect on attraction. To address the concern that the observed effects of the *Flatter* condition in Study 1 may have been due to novelty aversion, the difference in the email length, or the mere mention of a flatter hierarchy, we included a *Taller* condition in Study 2. For this condition, the second section mentioned that the firm has a “tall organizational structure” with “more levels of management than similarly sized startups in our industry.” To ensure that we successfully manipulated perceptions of hierarchy, the *Flatter* and *Taller* conditions included illustrative graphics.

Because the effects of an organization’s hierarchy may depend on the gender of its leadership, we test this boundary condition by manipulating the founder/CEO’s gender. In the second section (i.e.,

“About Us” in Figure A2), this gender manipulation was implemented by including a gender-typical name (i.e., “Jessica Chandler” or “Michael Chandler”), along with a gender-typical photo (for details, see Figure A3). Overall, our survey experiment consists of nine interventions because the organization’s hierarchical structure (i.e., no information vs. flatter vs. taller) and the founder/CEO’s gender (i.e., no information vs. woman vs. man) each have three conditions.

#### 4.5.3 Procedure

For this survey experiment, subjects were asked to first review the recruiting material, which contained the experimental manipulations.<sup>20</sup> Next, they advanced to a series of survey questions. The first set of questions checked whether the two treatments succeeded in manipulating the theoretical constructs of interest. For the organization’s hierarchical structure, subjects were asked to indicate how many management levels they believed the firm had between its founder/CEO and entry-level employees. In turn, for the founder/CEO’s gender, they were asked to specify their perception of the founder/CEO’s gender, the direct supervisor’s gender, and the proportion of women employees. The second set of questions asked about the extent to which the subjects found the firm attractive and their perceptions of various organizational characteristics (e.g., autonomy, informality, fit, workload, competition, and career advancement opportunities within the firm). This set also included a free-response question in which they shared their impression of a flatter hierarchy (if in either the *Flatter* or *No Information* condition) or a taller hierarchy (if in the *Taller* condition). The last set of questions inquired regarding the subject’s demographics.

#### 4.5.4 Measurement

**Dependent variable:** We measure attraction using Highhouse et al.’s (2003) five-item index of “employee recruitment and organization choice” (for details, see Appendix A9). We used this survey-based measure, rather than actual job applications, because Amazon prohibits recruiting its



workers for opportunities outside of MTurk. Subjects responded to each of the five items on a five-point Likert scale, ranging from “Strongly Disagree” (1 point) to “Strongly Agree” (5 points). After reverse-coding the second item in the index, we averaged these scores to calculate attraction.

**Independent variables:** Unlike Study 1, which had two experimental conditions, the survey experiment had three conditions: *Flatter*, *No Information*, and *Taller*. Setting the *Taller* condition as the comparison group, we coded these conditions using two binary variables indicating whether subjects received the *No Information* or *Flatter* conditions.

**Individual attributes:** We used the subjects’ responses to the demographic questions to record various individual attributes. These attributes include the subjects’ gender, current employment status, job title, work experience, and highest level of education. Among these attributes, the key variable of interest is *Woman*, which equals 1 if the subject identified as a woman, 0, otherwise.

**Mechanisms:** We measured subjects’ perceptions regarding a variety of the hiring firm’s characteristics using five-point Likert scales, including the *opportunity* for increased pay and promotion, the level of *autonomy* offered to employees,<sup>21</sup> the degree of *informality* of the firm’s procedures, and the procedural *fairness*. We also asked subjects how they perceived their *fit* with the organization, how much *workload* they would expect, whether the work environment would be *competitive* among coworkers, and the extent to which they believed the company would *succeed*. For details on how we measured each of these variables, see Appendix A10.

**Boundary condition:** To examine the potential boundary condition of leadership gender, we manipulated the founder/CEO’s gender and coded the variable *Woman Founder/CEO* as 1 if the subject received the woman founder/CEO condition, 0, otherwise.

#### 4.5.5 Summary statistics and randomization balance

In Table 3, the first column presents the summary statistics of subjects. Among the 9,000 subjects

recruited, 8,498 completed the survey and passed the attention check using the Founder/CEO's gender. The majority (58.4%) of these 8,498 subjects were women. In terms of the highest level of education, 44.0% completed less than a bachelor's degree, while 38.6% had only a bachelor's degree and 17.4% had a graduate degree. While most had more than six years of work experience (66.4%), only 24.7% held a managerial position in their most recent job. Overall, compared to the Study 1 participants, our survey experiment subjects were more likely to be women, less educated, and less likely to hold managerial positions. Compared to the population of U.S. workers, these subjects are more likely to be women, more educated, and more likely to be employed (for more information on their demographics, see Appendix A6).

Next, we move on to the second through fourth columns in Table 3, which present the summary statistics by treatment condition. The differences across these conditions are small and have 95% confidence intervals that substantially overlap with zero, implying that randomization was successful and that regression estimates can be interpreted as causal.

#### 4.5.6 Estimation

First, to check whether our manipulations worked as intended, we apply the equation:

$$Perc. Flatness_i = \gamma_0 + \gamma_1 No Information_i + \gamma_2 Flatter_i + \theta_i + \epsilon_i \quad (3)$$

where *Perc. Flatness<sub>i</sub>* stands for the reverse-coded, perceived number of hierarchical levels. The term  $\theta_i$  represents dummies for the founder/CEO's gender.

We then estimate the main effect of hierarchy on attraction using the specification:

$$Attraction_i = \beta_0 + \beta_1 No Information_i + \beta_2 Flatter_i + \theta_i + \epsilon_i \quad (4)$$

where  $\beta_1$  and  $\beta_2$  each represent the effect of the *No Information* and *Flatter* conditions relative to the *Taller* condition. Again,  $\theta_i$  represents dummies for the founder/CEO's gender.

Lastly, we examine whether the effect of hierarchy on attraction varies by gender using the

following equation

$$\begin{aligned} \text{Attraction}_i = & \beta_0 + \beta_1 \text{No Information}_i + \beta_2 \text{Flatter}_i + \beta_3 \text{Woman}_i \\ & + \beta_4 \text{No Information}_i \times \text{Woman}_i + \beta_5 \text{Flatter}_i \times \text{Woman}_i + \theta_i + \varepsilon_i \end{aligned} \quad (5)$$

where  $\beta_4$  denotes the extent to which attraction to the *No Information* condition relative to the *Taller* condition differs by gender, and  $\beta_5$  represents the extent to which attraction to the *Flatter* condition compared to the *Taller* condition varies by gender. The other terms are the same as in Equation 4. Accordingly, for each mechanism, we use its respective measure as the dependent variable (instead of *Attraction*) and apply the above estimation procedures. In addition, we supplement these OLS regressions with ANOVA tests.

As supplementary, non-preregistered analyses, we apply a two-stage least squares (2SLS) approach, where the randomized interventions (“Taller,” “No Information,” “Flatter”) are instrumental variables for our key explanatory variable of interest (the number of levels of hierarchy that a subject perceived). The second-stage regression provides a more precise, continuous operationalization of our theoretical construct of interest (i.e., perceived flatness), while maintaining the internal validity that comes from the random assignment. Whereas our field experimental design, like nearly all field experiments, could not account for subjects for whom the randomized interventions did not induce variation in the perceived level of hierarchical levels, the instrumental variables approach allows us to estimate the local average treatment effect (LATE), which identifies the effect among the subjects for whom the randomized intervention did induce variation in this explanatory variable (Angrist et al. 1996). To measure the main effect of hierarchy on attraction, we estimate the following 2SLS regressions, where the first stage is equivalent to Equation 3 and the second stage is:

$$\text{Attraction}_i = \beta_0 + \beta_1 \overbrace{\text{Perc. Flatness}_i} + \theta_i + \varepsilon_i \quad (6)$$

For the heterogenous effects by gender, we instead estimate the below second-stage equation:

$$Attraction_i = \beta_0 + \beta_1 \widehat{Perc. Flatness}_i + \beta_2 Woman_i + \beta_3 \widehat{Perc. Flatness}_i \times Woman_i + \theta_i + \varepsilon_i \quad (7)$$

where  $\widehat{Perc. Flatness}_i$  and  $\widehat{Perc. Flatness}_i \times Woman_i$  are the predicted values of perceived flatness and its interaction with the gender of the subject, respectively. Here, the term  $\beta_3$  represents the heterogeneous local average treatment effect by gender of the perceived flatness for those subjects that the treatments caused to perceive more or fewer hierarchical levels.

#### 4.5.7 Main results

The main results are reported in Table 4 and graphically illustrated using bar graphs in Figure 3. First, Model 1 demonstrates that our manipulation worked as intended (as also shown in panel (a) of Figure 3). A one-way ANOVA analysis shows significantly different effects between the conditions ( $F(2, 8495) = 4274.91, p = .000$ ). Compared to subjects randomly assigned to the *No Information* condition, those in the *Flatter* condition perceived the employer to have 1.3 fewer managerial levels ( $-1.303 \pm .027, p = .000$ ). Relative to the *Taller* condition, subjects assigned to the *No Information* condition perceived the firm to have 1.2 fewer managerial levels ( $-1.199 \pm .027, p = .000$ ). Finally, subjects assigned to the *Flatter* condition perceived the company to have 2.5 fewer levels than those in the *Taller* condition ( $-2.501 \pm .027, p = .000$ ).

Next, Model 2, which estimates Equation 4, indicates that subjects are, on average, least attracted to the *Taller* condition and most attracted to the *Flatter* condition ( $p = .000$ , for graphical representation, see panel (b) of Figure 3). Likewise, an ANOVA analysis reveals a significant difference between the three groups ( $F(2, 8495) = 59.73, p = .000$ ). A Tukey post-hoc test shows that participants were significantly more attracted to the *No Information* condition than the *Taller* condition ( $.151 \pm .025, p = .000$ ), to the *Flatter* condition than the *Taller* condition ( $.272 \pm .025,$

$p = .000$ ), and to the *Flatter* condition than the *No Information* condition ( $.221 \pm .025$ ,  $p = .000$ ). These results cast doubt on the alternative explanation for our Study 1 results that the observed effects of the *Flatter* condition, relative to the *No Information* condition, were driven by novelty aversion, the length of the email, or the mere mention of a flatter hierarchy. However, these results differ from Study 1's results in that they show a statistically significant main effect of a flatter hierarchy on attraction, thus supporting the argument that job seekers, on average, are more attracted to flatter hierarchies. One explanation for this difference is that, given their selection into remote work, MTurkers likely have a much higher baseline preference than our field-study subjects for the autonomy, flexibility, and informality associated with flatter hierarchies.

Model 3, which estimates Equation 5, tests whether men and women are differentially attracted to *flatter* and *taller* hierarchies. The results are graphically displayed in panel (c) of Figure 3. We find that, relative to men, women are somewhat less attracted to the *No Information* condition than to the *Taller* condition ( $p = .111$ ). However, compared to men, women are significantly less attracted to the *Flatter* condition than to the *Taller* condition ( $p = .0178$ ). Similarly, a two-way ANOVA analysis finds a significant interaction between the subject's gender and the firm's hierarchy on attraction ( $F(2, 8, 490) = 3.14$ ,  $p = .043$ ).

We now move on to Models 4 and 5 of Table 4, which show the results of the 2SLS approach. For both models, the first stage is strong, with a Kleibergen-Paap rk Wald F statistic of 5,273.5 and 1,622.6, respectively. Because we instrumented the perceived flatness (*Perc. Flatness*) using the randomized variables (i.e., *NoInformation* and *Flatter*) in the first-stage regression, the coefficient estimate of the variable *Perc. Flatness* in the second-stage regression can be interpreted as the causal effect of one fewer hierarchical level on attraction. Model 4 is consistent with Model 1 in that as subjects' perceptions of hierarchy increase, their attraction also increases

( $p < .001$ ). Model 5 indicates that, compared to men, women are disproportionately less attracted to a flatter hierarchy ( $p = .021$ ). Taken together, these main results of our survey experiment are consistent with our field-study findings that a flatter hierarchy reduces the female representation in the applicant pool. The key identifying assumption for the above 2SLS approach is the “exclusion restriction,”

or the requirement that the randomized intervention does not affect attraction except through subjects’ *perceived flatness*. Here, it is likely that characterizing an employer as flatter also induces variation in subjects’ perceptions of other employer traits (e.g., informality, autonomy, career advancement opportunities). As outlined in Section 2, however, because these perceptions follow from perceptions of hierarchy, they are not violations of the exclusion restriction, but rather the possible mechanisms whereby perceptions of hierarchy shape propensity to seek employment. We consider these mechanisms next.

#### 4.5.8 Post-hoc mechanism analyses

First, we examine how different hierarchical structures affect the average subject’s perceptions of a wide range of organizational characteristics. Then, we assess how these perceptions vary by gender. These results are presented in Table A9 in Appendix A12. Given the extensive list of organizational characteristics, we summarize the key findings here.

In line with the positive perspective on flatter hierarchies discussed in Section 2.3, subjects, on average, perceive these organizations as granting more *autonomy*. Furthermore, they anticipate that they will *fit* better in flatter employers and that these firms are more likely to *succeed*. However, subjects expect flatter hierarchies to be more *informal*, to provide less *opportunity* for career advancement, and to have less *competition* among employees. They do not, on average, anticipate significantly more or less *fairness* and *workload* in flatter organizations than in taller

ones. We find that women and men vary in some of these perceptions. Contrary to the positive view of flatter hierarchies discussed in Section 2.4, women subjects, compared to men, perceive flatter employers to offer them less *autonomy*. Furthermore, consistent with the negative perspective, women, compared to men, perceive flatter hierarchies as providing them significantly less *opportunity* for career progression, saddling them with a disproportionately heavier *workload*, and being more difficult for them to *fit* into. Yet, it does not appear that women are less attracted to flatter firms because they perceive more *competition* among employees in flatter organizations. Also, they do not differ significantly from men in whether they expect these firms to be more or less *informal*, *fair*, or likely to *succeed*.

These post-hoc analyses do not allow us to precisely isolate which mechanism drives women's relative aversion to flatter structures or, more generally, to conclude that women are not attracted to certain aspects of flatter hierarchies. However, they suggest a story that is broadly consistent with the negative perspective on flatter employers we discuss in Section 2.4. Relative to men, women job seekers are disproportionately less attracted to flatter employers because they perceive them as offering relatively less autonomy than men as well as engendering informal workplaces where they will encounter greater difficulty in fitting in and in progressing their careers.

#### **4.6 Discussion**

Flatter hierarchies have garnered much attention as “egalitarian” alternatives to taller, “authoritarian” hierarchies. Despite their potential benefits in motivating current employees, we present evidence that flatter hierarchies have a hidden cost of decreasing gender diversity in the applicant pool. Using a field experiment, we show that a flatter hierarchy does not have a significant effect on the size of the applicant pool but decreases its representation of women. Using a follow-up survey experiment, we replicate this latter finding and show that women's relative aversion to flatter

hierarchies corresponds to cross-gender differences in perceptions of career opportunities, fit, and workload. These findings have several implications for research on labor market gender segregation, organizational design, entrepreneurship, and human capital.

#### 4.6.1 Theoretical contributions

First, our work adds to the literature on labor market gender segregation. In particular, we extend prior work that has examined the application stage, showing that, relative to men, women may be more inclined to seek employment at organizations with particular characteristics (Abraham and Burbano 2022, Wiswall and Zafar 2018). We extend this work by examining the role of formal hierarchy—a fundamental feature of all organizations (Blau and Scott 2003, Burton and Obel 2004, Puranam 2018, Simon 1997)—in generating applicant sorting by gender. Drawing upon prior work on gendered organizations and organizational design, we develop opposing perspectives on how a firm’s flatter hierarchy impacts women and men job seekers’ perception and their propensity to apply to the firm. We show that, compared to taller structures, flatter hierarchies can exacerbate gender segregation by attracting disproportionately fewer women applicants. We find that this decrease in women’s representation reflects cross-gender differences in perceptions of career opportunities, fit, and workload. By considering how women and men differ in their perceptions of organizational characteristics (i.e., formal hierarchy) related to their job-seeking behavior, we answer recent calls to integrate “supply-” and “demand-side” perspectives on gender segregation in the labor market (Barbulescu and Bidwell 2013, Fernandez-Mateo and Kaplan 2018).

Second, our research contributes to the literature on organizational design by unveiling the role of organizational structure in attracting prospective employees. To date, this literature has primarily focused on how hierarchical structure motivates and coordinates existing employees (Burton and Obel 2004, Oldham and Hackman 1981, Puranam 2018), neglecting its effects on attracting new



employees. Although a flatter hierarchy may benefit firms by motivating current employees and spurring their creativity (Keum and See 2017, Reitzig 2022), we find that it has an unintended consequence of decreasing the diversity among prospective employees. Because the likelihood of an employer hiring employees with a certain attribute is typically proportional to their representation in its applicant pool (Fernandez and Abraham 2011, Fernandez-Mateo and Fernandez 2016, Petersen et al. 2000), this decrease in applicant diversity may negatively impact employee diversity within the organization.

Third, we advance entrepreneurship research by uncovering an unintended, potentially long-lasting consequence of a flatter hierarchy for startups. Past research has argued that new ventures should be “organic” with a flatter hierarchy to quickly and flexibly adapt to their turbulent environment (Burns and Stalker 1961). Adding to the burgeoning stream of work reevaluating this received wisdom (e.g., Foss and Klein 2022, Keum and See 2017, Lee 2022, Lee and Csaszar 2020, Reitzig 2022), our study reveals that a flatter hierarchy can have an unintended consequence of reducing the gender diversity in the applicant pool and thus aggravating the problem of “diversity debt” (Wu 2017). This finding demonstrates how initial decisions that founders make regarding organizational structure may exert a lasting influence on the firm’s evolution and success (Alexy et al. 2021, Baron et al. 1996, 2007, Beckman and Burton 2008).

Lastly, this study extends work on human capital by providing an empirical examination of the endogenous process through which individuals self-select into firms. Although extant studies have enumerated various firm characteristics that affect this selection process (e.g., leadership gender, firm size, organizational status, corporate social responsibility, Abraham and Burbano 2022, Bidwell et al. 2015, Burbano 2016, Campero and Kacperczyk 2020, Elfenbein et al. 2010), they have yet to examine the effects of formal hierarchy, which is “one of the defining features of formal organizations”

(Sørensen and Sharkey 2014:329). Given the increasing trend of companies highlighting their flatter hierarchies in their recruiting efforts, we provide the first study that assesses how perceptions of a firm's hierarchy shape individuals' self-selection into firms. We also explore mechanisms underlying this relationship by examining how women and men differ in their perceptions. By doing so, we shed light on the selection process that determines founding team characteristics (Shah et al. 2019).

#### 4.6.2 Limitations and future research

Like all research, this study has limitations. First, although our field experiment may have more external validity than typical lab experiments, it involved a single firm in a specific industry (i.e., a healthcare startup). Moreover, this experiment considered just two job positions (business development representative and software engineer), and its subjects were limited to job seekers with at least a bachelor's degree. Although we find little evidence of heterogeneous effects across jobs and by individual attributes and our field-study findings are largely consistent in our follow-up survey experiment using a broader population, future studies could enhance the generalizability of our results or delineate their boundary conditions. For example, by employing two-dimensional randomized treatments (i.e., by randomizing not just the employer's hierarchy but also other characteristics), these studies could evaluate whether job seekers' attraction to flatter hierarchies varies across hierarchical ranks (i.e., whether job seekers currently hold a managerial or entry-level position, whether the hiring firm's open position is a managerial or entry-level job), occupations (e.g., gender composition, task modularity, coordination requirement, knowledge intensity), firms (e.g., firm age, firm size, organizational culture, incentive system, mentorship), industries (e.g., high-tech, industry life-cycle), geographies, and time.

Second, our study examined only one dimension of organizational structure as the explanatory variable of interest. Future research could extend this study by probing into other structural

dimensions (e.g., the horizontal division of tasks), work arrangements (e.g., remote work), or human resources policies (e.g., job training) that, like formal hierarchy, do not explicitly evoke gender or diversity.

Third, in accordance with prior studies on gendered organizations (e.g., Acker 1990, Ferguson 1984, Nicholson 2010, Reskin 1988), our study focused on binary gender and measured its composition as the outcome. Although our theoretical arguments and empirical findings might broadly apply to gender non-conforming job seekers given their historically marginalized status, it was empirically infeasible to distinguish these individuals in our field experiment because ZipRecruiter does not provide information on gender and Gender API's algorithm cannot identify these subjects using only their names. In addition, although our survey experiment included a question that captures non-binary gender identities, such identities were extremely rare (i.e., less than 1%) in the sample, making it difficult to estimate the heterogeneous effects for these individuals with a meaningful degree of precision. Building upon this study, future work may develop theories and employ alternative sampling strategies to elucidate how gender non-conforming job seekers perceive flatter hierarchies. Fourth, in addition to job seekers in the LGBTQ+ community, future research may explore other historically marginalized groups (e.g., racial minorities and immigrants). Although members of these groups may not encounter the same gender-specific expectations that women face, they may share women's perceptions that flatter organizations leave room for the dominant group to consolidate power in ways that inhibit their career progression. Supporting this argument, several practitioners have suggested that these organizations disadvantage not only women but also other minority groups (e.g., Diamond 2019, Finley 2014, Flower Horne 2014, Hunt 2017, Mont 2017). If this is the case, flatter hierarchies may also substantially reduce employee diversity in terms of other identities than gender.

Fifth, we surveyed an extensive yet finite list of potential mechanisms underlying how a firm's formal hierarchy may shape women and men job seeker's perceptions and their attraction to the firm. Future studies could complement our research by exploring other possible mechanisms that may explain why women and men react differently to flatter or taller hierarchies. These studies may provide additional insights into how flatter organizations could mitigate the under-representation of women in their applicant pool.

Finally, our study centered only on employee attraction. Yet, it remains unclear whether a flatter hierarchy is more or less likely than a taller one to retain employees, particularly women. If women job seekers' relatively negative perceptions of flatter hierarchies are accurate, flatter organizations may also encounter difficulty in retaining women. Instead, if these perceptions are inaccurate and flatter hierarchies are more advantageous for women, flatter employers may more successfully retain them. The women-retaining advantages of flatter structures might represent a countervailing force that offsets their diversity-reducing effects on the applicant pool. Future work examining retention will thus elucidate to what extent flatter hierarchies augment or ameliorate workplace gender diversity.

#### **4.7 Conclusion**

In sum, our study shows firms that tout flatter hierarchies attract fewer women to their applicant pools. Yet, many important insights remain to be uncovered regarding the relationship between organizational structure and human capital acquisition. We hope our study serves as a foundation for future exploration in this area.

## 4.8 Works Cited

- Abraham, Mabel. "Gender-Role Incongruity and Audience-Based Gender Bias: An Examination of Networking among Entrepreneurs." *Administrative Science Quarterly* 65, no. 1 (March 1, 2020): 151–80. <https://doi.org/10.1177/0001839219832813>.
- . "Pay Formalization Revisited: Considering the Effects of Manager Gender and Discretion on Closing the Gender Wage Gap." *Academy of Management Journal* 60, no. 1 (February 2017): 29–54. <https://doi.org/10.5465/amj.2013.1060>.
- Abraham, Mabel, and Vanessa Burbano. "Congruence Between Leadership Gender and Organizational Claims Affects the Gender Composition of the Applicant Pool: Field Experimental Evidence." *Organization Science*, March 17, 2021, orsc.2021.1442. <https://doi.org/10.1287/orsc.2021.1442>.
- Acker, Joan. "Hierarchies, Jobs, Bodies:: A Theory of Gendered Organizations." *Gender & Society* 4, no. 2 (June 1, 1990): 139–58. <https://doi.org/10.1177/089124390004002002>.
- Alexy, Oliver, Katharina Poetz, Phanish Puranam, and Markus Reitzig. "Adaptation or Persistence? Emergence and Revision of Organization Designs in New Ventures." *Organization Science* 32, no. 6 (November 2021): 1439–72. <https://doi.org/10.1287/orsc.2021.1431>.
- Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91, no. 434 (1996): 444–55. <https://doi.org/10.2307/2291629>.
- Angrist, Joshua D., and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2008.
- Babcock, Linda, and Sara Laschever. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press, 2009.
- Barber, Alison E., and Mark V. Roehling. "Job Postings and the Decision to Interview: A Verbal Protocol Analysis." *Journal of Applied Psychology* 78, no. 5 (October 1993): 845–56. <https://doi.org/10.1037/0021-9010.78.5.845>.
- Barbulescu, Roxana, and Matthew Bidwell. "Do Women Choose Different Jobs from Men? Mechanisms of Application Segregation in the Market for Managerial Workers." *Organization Science* 24, no. 3 (June 2013): 737–56. <https://doi.org/10.1287/orsc.1120.0757>.
- Baron, James N. "Organizational Perspectives on Stratification." *Annual Review of Sociology* 10 (1984): 37–69.
- Baron, James N., M Diane Burton, and Michael T Hannan. "The Road Taken: Origins and Evolution of Employment Systems in Emerging Companies," n.d.

- Baron, James N., Alison Davis-Blake, and William T. Bielby. “The Structure of Opportunity: How Promotion Ladders Vary Within and Among Organizations.” *Administrative Science Quarterly* 31, no. 2 (1986): 248–73. <https://doi.org/10.2307/2392790>.
- Baron, James N., Michael T. Hannan, Greta Hsu, and Özgecan Koçak. “In the Company of Women: Gender Inequality and the Logic of Bureaucracy in Start-Up Firms.” *Work and Occupations* 34, no. 1 (February 1, 2007): 35–66. <https://doi.org/10.1177/0730888406296945>.
- Barrymore, Nathan, Cristian L. Dezső, and Benjamin C. King. “Gender and Competitiveness When Earning for Others: Experimental Evidence and Implications for Sponsorship.” *Strategic Management Journal* 43, no. 5 (2022): 905–34. <https://doi.org/10.1002/smj.3353>.
- Bartling, Björn, Ernst Fehr, and Holger Herz. “The Intrinsic Value of Decision Rights.” *Econometrica* 82, no. 6 (2014): 2005–39. <https://doi.org/10.3982/ECTA11573>.
- Beckman, Christine M., and M. Diane Burton. “Founding the Future: Path Dependence in the Evolution of Top Management Teams from Founding to IPO.” *Organization Science* 19, no. 1 (2008): 3–24.
- Bennett, Victor Manuel. “Organization and Bargaining: Sales Process Choice at Auto Dealerships.” *Management Science* 59, no. 9 (September 2013): 2003–18. <https://doi.org/10.1287/mnsc.1120.1691>.
- Bidwell, Matthew, Shinjae Won, Roxana Barbulescu, and Ethan Mollick. “I Used to Work at Goldman Sachs! How Firms Benefit from Organizational Status in the Market for Human Capital.” *Strategic Management Journal* 36, no. 8 (2015): 1164–73. <https://doi.org/10.1002/smj.2272>.
- Bielby, William T., and James N. Baron. “Men and Women at Work: Sex Segregation and Statistical Discrimination.” *American Journal of Sociology* 91, no. 4 (1986): 759–99.
- Blau, Francine D., and Wallace E. Hendricks. “Occupational Segregation by Sex: Trends and Prospects.” *The Journal of Human Resources* 14, no. 2 (1979): 197–210. <https://doi.org/10.2307/145642>.
- Blau, Francine D., Patricia Simpson, and Deborah Anderson. “Continuing Progress? Trends in Occupational Segregation in the United States over the 1970s and 1980s.” *Feminist Economics* 4, no. 3 (January 1, 1998): 29–71. <https://doi.org/10.1080/135457098338301>.
- Blau, Peter Michael, and W. Richard Scott. *Formal Organizations: A Comparative Approach*. Stanford University Press, 2003.
- Brockner, Joel, and Laury Adsit. “The Moderating Impact of Sex on the Equity–Satisfaction Relationship: A Field Study.” *Journal of Applied Psychology* 71, no. 4 (November 1986): 585–90. <https://doi.org/10.1037/0021-9010.71.4.585>.

- Burbano, Vanessa C. “Social Responsibility Messages and Worker Wage Requirements: Field Experimental Evidence from Online Labor Marketplaces.” *Organization Science* 27, no. 4 (August 2016): 1010–28. <https://doi.org/10.1287/orsc.2016.1066>.
- Graicunas, V.A. “Relationship in Organization.” In *Papers on the Science of Administration*, edited by L Gulick and L. Urwick, 183–87. New York: Institute of Public Administration, 1937.
- Gruenfeld, Deborah H., and Larissa Z. Tiedens. “Organizational Preferences and Their Consequences.” In *Handbook of Social Psychology, Vol. 2, 5th Ed*, 1252–87. Hoboken, NJ, US: John Wiley & Sons, Inc., 2010. <https://doi.org/10.1002/9780470561119.socpsy002033>.
- Halac, Marina, and Andrea Prat. “Managerial Attention and Worker Performance.” *American Economic Review* 106, no. 10 (October 2016): 3104–32. <https://doi.org/10.1257/aer.20140772>.
- Halevy, Nir, Eileen Y. Chou, and Adam D. Galinsky. “A Functional Model of Hierarchy: Why, How, and When Vertical Differentiation Enhances Group Performance.” *Organizational Psychology Review* 1, no. 1 (February 1, 2011): 32–52. <https://doi.org/10.1177/2041386610380991>.
- Hamburger, Liz. “The Flat Hierarchy Fallacy.” *Medium* (blog), May 4, 2021. <https://lizhamburger.medium.com/the-flat-hierarchy-fallacy-ce8b4c3dfc85>.
- He, Haoran, David Neumark, and Qian Weng. “Do Workers Value Flexible Jobs? A Field Experiment.” *Journal of Labor Economics* 39, no. 3 (July 2021): 709–38. <https://doi.org/10.1086/711226>.
- Highhouse, Scott, Filip Lievens, and Evan F. Sinar. “Measuring Attraction to Organizations.” *Educational and Psychological Measurement* 63, no. 6 (December 1, 2003): 986–1001. <https://doi.org/10.1177/0013164403258403>.
- Hunt, Bill. “The Myth of the Flat Organization.” *Medium* (blog), June 13, 2017. <https://krusynth.medium.com/the-myth-of-the-flat-organization-b522529bc24a>.
- Kahn, Arnold, Judith E. Krulewitz, Virginia E. O’Leary, and Helmut Lamm. “Equity and Equality: Male and Female Means to a Just End.” *Basic and Applied Social Psychology* 1, no. 2 (June 1, 1980): 173–97. [https://doi.org/10.1207/s15324834basp0102\\_6](https://doi.org/10.1207/s15324834basp0102_6).
- Kalev, Alexandra. “Cracking the Glass Cages? Restructuring and Ascriptive Inequality at Work.” *American Journal of Sociology* 114, no. 6 (May 2009): 1591–1643. <https://doi.org/10.1086/597175>.
- Kanter, Rosabeth Moss. *Men and Women of the Corporation: New Edition*. Basic Books, 2008.
- Kees, Jeremy, Christopher Berry, Scot Burton, and Kim Sheehan. “An Analysis of Data Quality: Professional Panels, Student Subject Pools, and Amazon’s Mechanical Turk.” *Journal of*

- Advertising* 46, no. 1 (January 2, 2017): 141–55.  
<https://doi.org/10.1080/00913367.2016.1269304>.
- Keum, Dongil D., and Kelly E. See. “The Influence of Hierarchy on Idea Generation and Selection in the Innovation Process.” *Organization Science* 28, no. 4 (August 2017): 653–69. <https://doi.org/10.1287/orsc.2017.1142>.
- Kooti, Farshad, Luca Maria Aiello, Mihajlo Grbovic, Kristina Lerman, and Amin Mantrach. “Evolution of Conversations in the Age of Email Overload.” In *Proceedings of the 24th International Conference on World Wide Web*, 603–13. WWW ’15. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, 2015. <https://doi.org/10.1145/2736277.2741130>.
- Lambrecht, Anja, and Catherine Tucker. “Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads.” *Management Science* 65, no. 7 (July 2019): 2966–81. <https://doi.org/10.1287/mnsc.2018.3093>.
- Lee, Michael Y., and Amy C. Edmondson. “Self-Managing Organizations: Exploring the Limits of Less-Hierarchical Organizing.” *Research in Organizational Behavior* 37 (January 1, 2017): 35–58. <https://doi.org/10.1016/j.riob.2017.10.002>.
- Lee, Saerom, and Felipe A. Csaszar. “Cognitive and Structural Antecedents of Innovation: A Large-Sample Study.” *Strategy Science* 5, no. 2 (June 2020): 71–97. <https://doi.org/10.1287/stsc.2020.0107>.
- Lee, Saerom (Ronnie). “The Myth of the Flat Start-up: Reconsidering the Organizational Structure of Start-Ups.” *Strategic Management Journal* 43, no. 1 (2022): 58–92. <https://doi.org/10.1002/smj.3333>.
- Lee, Saerom (Ronnie), and J. Daniel Kim. “When Do Startups Scale? Large-Scale Evidence from Job Postings.” SSRN Scholarly Paper. Rochester, NY, August 1, 2023. <https://doi.org/10.2139/ssrn.4015530>.
- Lerchenmueller, Marc J., Olav Sorenson, and Anupam B. Jena. “Gender Differences in How Scientists Present the Importance of Their Research: Observational Study.” *BMJ* 367 (December 16, 2019): l6573. <https://doi.org/10.1136/bmj.l6573>.
- Lumpkin, G.T., Claudia C. Cogliser, and Dawn R. Schneider. “Understanding and Measuring Autonomy: An Entrepreneurial Orientation Perspective.” *Entrepreneurship Theory and Practice* 33, no. 1 (January 1, 2009): 47–69. <https://doi.org/10.1111/j.1540-6520.2008.00280.x>.
- March, James, ed. “Social Structure and Organizations.” In *Handbook of Organizations (RLE: Organizations)*. Routledge, 1965.
- Mas, Alexandre, and Amanda Pallais. “Valuing Alternative Work Arrangements.” *American Economic Review* 107, no. 12 (December 2017): 3722–59. <https://doi.org/10.1257/aer.20161500>.



- Mihaljević, Helena, Ivana Müller, Katja Dill, and Aysel Yollu-Tok. “Towards Gender-Inclusive Job Postings: A Data-Driven Comparison of Augmented Writing Technologies.” *PLOS ONE* 17, no. 9 (September 9, 2022): e0274312. <https://doi.org/10.1371/journal.pone.0274312>.
- Mont, Simon. “Autopsy of a Failed Holacracy: Lessons in Justice, Equity, and Self-Management.” Non Profit News | Nonprofit Quarterly, January 9, 2017. <https://nonprofitquarterly.org/autopsy-failed-holacracy-lessons-justice-equity-self-management/>.
- Moss, Aaron J., Cheskie Rosenzweig, Jonathan Robinson, Shalom N. Jaffe, and Leib Litman. “Is It Ethical to Use Mechanical Turk for Behavioral Research? Relevant Data from a Representative Survey of MTurk Participants and Wages.” *Behavior Research Methods*, May 22, 2023. <https://doi.org/10.3758/s13428-022-02005-0>.
- Murciano-Goroff, Raviv. “Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers.” *Management Science* 68, no. 5 (May 2022): 3262–81. <https://doi.org/10.1287/mnsc.2021.4077>.
- Murray, Henry Alexander. *Explorations in Personality: A Clinical and Experimental Study of Fifty Men of College Age*. Oxford University Press, 1938.
- Nelson, Robert L., and William P. Bridges. *Legalizing Gender Inequality: Courts, Markets and Unequal Pay for Women in America*. Cambridge University Press, 1999.
- Nicholson, Nigel. “Gender and the Future of the Hierarchical Organization.” *Harvard Business Review*, June 7, 2010. <https://hbr.org/2010/06/gender-and-the-future-of-the-o>.
- Niederle, Muriel, and Lise Vesterlund. “Gender and Competition.” *Annual Review of Economics* 3, no. 1 (2011): 601–30. <https://doi.org/10.1146/annurev-economics-111809-125122>.
- Oldham, Greg R., and J. Richard Hackman. “Relationships Between Organizational Structure and Employee Reactions: Comparing Alternative Frameworks.” *Administrative Science Quarterly* 26, no. 1 (1981): 66–83. <https://doi.org/10.2307/2392600>.
- Padavic, Irene, and Barbara F. Reskin. *Women and Men at Work*. Pine Forge Press, 2002.
- Perry, Elissa L., Alison Davis-Blake, and Carol T. Kulik. “Explaining Gender-Based Selection Decisions: A Synthesis of Contextual and Cognitive Approaches.” *The Academy of Management Review* 19, no. 4 (1994): 786–820. <https://doi.org/10.2307/258745>.
- Petersen, Trond, and Laurie A. Morgan. “Separate and Unequal: Occupation-Establishment Sex Segregation and the Gender Wage Gap.” *American Journal of Sociology* 101, no. 2 (1995): 329–65.
- Petersen, Trond, Ishak Saporta, and Marc-David L. Seidel. “Offering a Job: Meritocracy and Social Networks.” *American Journal of Sociology* 106, no. 3 (2000): 763–816. <https://doi.org/10.1086/318961>.

- Phillips, Damon J. “Organizational Genealogies and the Persistence of Gender Inequality: The Case of Silicon Valley Law Firms.” *Administrative Science Quarterly* 50, no. 3 (September 1, 2005): 440–72. <https://doi.org/10.2189/asqu.2005.50.3.440>.
- Puranam, Phanish. *The Microstructure of Organizations*. Oxford University Press, 2018. <https://doi.org/10.1093/oso/9780199672363.001.0001>.
- Puranam, Phanish, and Dorthe Døjbak Håkonsson. “Valve’s Way.” *Journal of Organization Design* 4, no. 2 (June 23, 2015): 2–4. <https://doi.org/10.7146/jod.20152>.
- Reingold, Jennifer. “How a Radical Shift Left Zappos Reeling.” *Fortune*. Accessed December 5, 2023. <https://fortune.com/longform/zappos-tony-hsieh-holacracy/>.
- Reitzig, Markus. *Get Better at Flatter: A Guide to Shaping and Leading Organizations with Less Hierarchy*. Springer International Publishing, 2022.
- Reitzig, Markus, and Olav Sorenson. “Biases in the Selection Stage of Bottom-up Strategy Formulation.” *Strategic Management Journal* 34, no. 7 (2013): 782–99. <https://doi.org/10.1002/smj.2047>.
- Reskin, Barbara. “Sex Segregation in the Workplace.” *Annual Review of Sociology* 19 (1993): 241–70.
- Reskin, Barbara F. “Bringing the Men Back in: Sex Differentiation and the Devaluation of Women’s Work.” *Gender and Society* 2, no. 1 (1988): 58–81.
- Reskin, Barbara F., Patricia A. Roos, Katharine M. Donato, Polly A. Phipps, Barbara J. Thomas, Chloe E. Bird, Linda A. Detman, and Thomas Steiger. *Job Queues, Gender Queues: Explaining Women’s Inroads into Male Occupations*. Temple University Press, 1990. <https://www.jstor.org/stable/j.ctt14btcj8>.
- Ridgeway, Cecilia L. *Framed by Gender: How Gender Inequality Persists in the Modern World*. Oxford University Press, USA, 2011.
- Saxenian, AnnaLee. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128, With a New Preface by the Author*. Harvard University Press, 1996.
- Sebo, Paul. “Performance of Gender Detection Tools: A Comparative Study of Name-to-Gender Inference Services.” *Journal of the Medical Library Association* 109, no. 3 (October 5, 2021): 414–21. <https://doi.org/10.5195/jmla.2021.1185>.
- Shah, Sonali K., Rajshree Agarwal, and Raj Echambadi. “Jewels in the Crown: Exploring the Motivations and Team Building Processes of Employee Entrepreneurs.” *Strategic Management Journal* 40, no. 9 (2019): 1417–52. <https://doi.org/10.1002/smj.3027>.
- Shaver, J. Myles. “Interpreting Interactions in Linear Fixed-Effect Regression Models: When Fixed-Effect Estimates Are No Longer Within-Effects.” *Strategy Science* 4, no. 1 (March 2019): 25–40. <https://doi.org/10.1287/stsc.2018.0065>.

- Siegel, Jordan I., Amir N. Licht, and Shalom H. Schwartz. "Egalitarianism, Cultural Distance, and Foreign Direct Investment: A New Approach." *Organization Science* 24, no. 4 (2013): 1174–94.
- Simon, Herbert Alexander. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*. Free Press, 1976.
- Slough, Tara. "Phantom Counterfactuals." *American Journal of Political Science* 67, no. 1 (2023): 137–53. <https://doi.org/10.1111/ajps.12715>.
- Snowberg, Erik, and Leeat Yariv. "Testing the Waters: Behavior across Participant Pools." *American Economic Review* 111, no. 2 (February 2021): 687–719. <https://doi.org/10.1257/aer.20181065>.
- Sørensen, Jesper B., and Amanda J. Sharkey. "Entrepreneurship as a Mobility Process." *American Sociological Review* 79, no. 2 (April 1, 2014): 328–49. <https://doi.org/10.1177/0003122414521810>.
- Sorenson, Olav. "Flat Firms, Complementary Choices, Employee Effort, and the Pyramid Principle." *Journal of Organization Design* 11, no. 1 (March 1, 2022): 11–14. <https://doi.org/10.1007/s41469-022-00115-9>.
- Spector, Paul E. "Perceived Control by Employees: A Meta-Analysis of Studies Concerning Autonomy and Participation at Work." *Human Relations* 39, no. 11 (November 1, 1986): 1005–16. <https://doi.org/10.1177/001872678603901104>.
- Stern, Scott. "Do Scientists Pay to Be Scientists?" *Management Science* 50, no. 6 (2004): 835–53.
- Tetlock, Philip E. "Accountability: The Neglected Social Context of Judgment and Choice." *Research in Organizational Behavior* 7 (1985): 297–332.
- The Elephant in the Valley. "The Elephant in the Valley." Accessed December 5, 2023. <https://www.elephantinthevalley.com>.
- Tiedens, Larissa Z., Miguel M. Unzueta, and Maia J. Young. "An Unconscious Desire for Hierarchy? The Motivated Perception of Dominance Complementarity in Task Partners." *Journal of Personality and Social Psychology* 93, no. 3 (September 2007): 402–14. <https://doi.org/10.1037/0022-3514.93.3.402>.
- Tomaskovic-Devey, Donald. *Gender & Racial Inequality at Work: The Sources and Consequences of Job Segregation*. Cornell University Press, 1993.
- Wiswall, Matthew, and Basit Zafar. "Preference for the Workplace, Investment in Human Capital, and Gender\*." *The Quarterly Journal of Economics* 133, no. 1 (February 1, 2018): 457–507. <https://doi.org/10.1093/qje/qjx035>.

Wu, Susan. “Welcome to Diversity Debt: The Crisis That Could Sink Uber.” *Wired*. Accessed December 5, 2023. <https://www.wired.com/2017/03/welcome-to-diversity-debt-the-crisis-that-could-sink-uber/>.

Yu, Siyu, Lindred L. Greer, Nir Halevy, and Lisanne van Bunderen. “On Ladders and Pyramids: Hierarchy’s Shape Determines Relationships and Performance in Groups.” *Personality and Social Psychology Bulletin* 45, no. 12 (December 1, 2019): 1717–33. <https://doi.org/10.1177/0146167219842867>.

## **4.9 Tables**

Table 4.9.1 Summary statistics for Study 1.

	All	By condition		By position	
		<i>No Information</i>	<i>Flatter</i>	<i>BD</i>	<i>SW</i>
<b>Outcomes</b>					
<i>Click</i>	0.2849 (0.4514)	0.2863 (0.4521)	0.2835 (0.4508)	0.1890 (0.3916)	0.3831 (0.4862)
<i>Apply</i>	0.0973 (0.2964)	0.0989 (0.2985)	0.0958 (0.2944)	0.0523 (0.2226)	0.1435 (0.3506)
<b>Hierarchy</b>					
<i>Flatter</i>	0.5009 (0.5000)	0 (0)	1 (0)	0.5005 (0.5001)	0.5014 (0.5001)
<b>Gender</b>					
<i>Woman</i>	0.2688 (0.4433)	0.2701 (0.4441)	0.2674 (0.4427)	0.3059 (0.4608)	0.2307 (0.4214)
<b>Interaction</b>					
<i>Flatter × Woman</i>	0.1340 (0.3406)	0 (0)	0.2674 (0.4427)	0.1539 (0.3609)	0.1135 (0.3173)
<b>Job position</b>					
<i>Software</i>	0.4941 (0.5000)	0.4936 (0.5000)	0.4945 (0.5000)	0 (0)	1 (0)
<b>Current job title</b>					
<i>Manager</i>	0.3178 (0.4657)	0.3188 (0.4661)	0.3168 (0.4653)	0.5253 (0.4994)	0.1053 (0.3069)
<b>Years of work experience</b>					
<i>Work Experience</i>	13.8536 (8.8420)	13.8297 (8.7179)	13.8776 (8.9655)	16.3175 (9.1617)	10.8187 (7.3770)
<b>Highest degree</b>					
<i>Bachelor's</i>	0.2712 (0.4446)	0.2634 (0.4405)	0.2790 (0.4486)	0.2980 (0.4574)	0.2438 (0.4295)
<i>Masters</i>	0.3578 (0.4794)	0.3594 (0.4799)	0.3561 (0.4789)	0.3173 (0.4655)	0.3993 (0.4898)
<i>Doctorate</i>	0.0173 (0.1303)	0.0174 (0.1308)	0.0171 (0.1297)	0.0186 (0.1352)	0.0159 (0.1250)
<b>Other</b>					
<i>Willing to Relocate</i>	0.3684 (0.4824)	0.3708 (0.4831)	0.3661 (0.4818)	0.3288 (0.4698)	0.4091 (0.4917)
No. observations	8,167	4,076	4,091	4,132	4,035

*Note.* Standard deviations in parentheses.

Table 4.9.2 Results using OLS regression in Study 1. Panel (a) employs the full sample, while panels (b) and (c) represent splits sample analysis by the job position that was being advertised.

(a) Full sample

Outcome:	Click					Apply				
	Model 1 (All)	Model 2 (All)	Model 3 (All)	Model 4 (Women only)	Model 5 (Men only)	Model 6 (All)	Model 7 (All)	Model 8 (All)	Model 9 (Women only)	Model 10 (Men only)
<i>Flatter</i>	-0.0029 (0.0098) [0.7639]	0.0103 (0.0115) [0.3678]	-0.0029 (0.0098) [0.7645]	-0.0390 (0.0186) [0.0362]	0.0103 (0.0115) [0.3682]	-0.0031 (0.0065) [0.6291]	0.0058 (0.0078) [0.4515]	-0.0032 (0.0065) [0.6261]	-0.0277 (0.0117) [0.0177]	0.0058 (0.0078) [0.4524]
<i>Woman</i>		0.0280 (0.0157) [0.0752]	0.0033 (0.0110) [0.7662]				0.0069 (0.0104) [0.5083]	-0.0099 (0.0070) [0.1578]		
<i>Flatter × Woman</i>		-0.0493 (0.0218) [0.0242]	-0.0504 (0.0219) [0.0215]				-0.0335 (0.0140) [0.0170]	-0.0358 (0.0140) [0.0108]		
Position FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	FE	FE	DDFE	FE	FE	FE	FE	DDFE	FE	FE
No. observations	8,167	8,167	8,167	2,195	5,972	8,167	8,167	8,167	2,195	5,972
R <sup>2</sup>	0.046	0.047	0.047	0.047	0.047	0.024	0.025	0.025	0.025	0.024

(b) Business development (BD)

Outcome:	Click				Apply			
	Model 11 (BD)	Model 12 (BD)	Model 13 (BD Women only)	Model 14 (BD Men only)	Model 15 (BD)	Model 16 (BD)	Model 17 (BD Women only)	Model 18 (BD Men only)
<i>Flatter</i>	0.0011 (0.0123) [0.9298]	0.0221 (0.0146) [0.1303]	-0.0468 (0.0224) [0.0368]	0.0221 (0.0146) [0.1302]	0.0076 (0.0069) [0.2698]	0.0162 (0.0084) [0.0553]	-0.0117 (0.0121) [0.3331]	0.0162 (0.0084) [0.0553]
<i>Woman</i>		0.0431 (0.0194) [0.0266]				0.0082 (0.0106) [0.4400]		
<i>Flatter × Woman</i>		-0.0690 (0.0268) [0.0100]				-0.0279 (0.0147) [0.0585]		
No. observations	4,132	4,132	1,264	2,868	4,132	4,132	1,264	2,868
R <sup>2</sup>	0.000	0.002	0.004	0.001	0.000	0.001	0.001	0.001

(c) Software (SW)

Outcome:	Click				Apply			
	Model 19 (SW)	Model 20 (SW)	Model 21 (SW Women only)	Model 22 (SW Men only)	Model 23 (SW)	Model 24 (SW)	Model 25 (SW Women only)	Model 26 (SW Men only)
<i>Flatter</i>	-0.0090 (0.0153) [0.5553]	-0.0020 (0.0175) [0.9112]	-0.0326 (0.0319) [0.3074]	-0.0020 (0.0175) [0.9112]	-0.0142 (0.0110) [0.1995]	-0.0037 (0.0127) [0.7681]	-0.0495 (0.0221) [0.0255]	-0.0037 (0.0127) [0.7680]
<i>Woman</i>		0.0172 (0.0258) [0.5046]				0.0077 (0.0190) [0.6875]		
<i>Flatter × Woman</i>		-0.0307 (0.0364) [0.3995]				-0.0457 (0.0255) [0.0730]		
No. observations	4,035	4,035	931	3,104	4,035	4,035	931	3,104
R <sup>2</sup>	0.000	0.000	0.001	0.000	0.000	0.001	0.005	0.000

Note. Robust standard errors in parentheses and p-values in brackets. Models 3 and 8 use double-demeaned fixed effects (DDFE) estimators (Giesecke and Schmidt-Catran 2022)

Table 4.9.3 Summary statistics in Study 2.

	All	By condition		
		Taller	No Information	Flatter
<b>Outcome</b>				
<i>Attraction</i>	3.6516 (0.9480)	3.5095 (1.0048)	3.6607 (0.8887)	3.7816 (0.9273)
<b>Hierarchy</b>				
<i>Taller</i>	0.3315 (0.4708)	1 (0)	0 (0)	0 (0)
<i>No Information</i>	0.3291 (0.4699)	0 (0)	1 (0)	0 (0)
<i>Flatter</i>	0.3394 (0.4735)	0 (0)	0 (0)	1 (0)
<b>Gender</b>				
<i>Woman</i>	0.5841 (0.4929)	0.5875 (0.4924)	0.5896 (0.4920)	0.5756 (0.4943)
<b>Interactions</b>				
<i>Taller × Woman</i>	0.1948 (0.3960)	0.5875 (0.4924)	0 (0)	0 (0)
<i>No Information × Woman</i>	0.1940 (0.3955)	0 (0)	0.5896 (0.4920)	0 (0)
<i>Flatter × Woman</i>	0.1953 (0.3965)	0 (0)	0 (0)	0.5756 (0.4943)
<b>Current job title</b>				
<i>Manager</i>	0.2475 (0.4316)	0.2588 (0.4380)	0.2481 (0.4320)	0.2358 (0.4246)
<b>Work experience</b>				
<i>No experience</i>	0.0595 (0.2367)	0.0618 (0.2408)	0.0597 (0.2370)	0.0572 (0.2323)
<i>1-3 years</i>	0.1344 (0.3411)	0.1313 (0.3378)	0.1237 (0.3293)	0.1477 (0.3549)
<i>4-6 years</i>	0.1419 (0.3490)	0.1402 (0.3473)	0.1416 (0.3487)	0.1439 (0.3510)
<i>7-9 years</i>	0.1098 (0.3126)	0.1154 (0.3195)	0.1065 (0.3086)	0.1075 (0.3098)
<i>10+ years</i>	0.5542 (0.4971)	0.5513 (0.4975)	0.5681 (0.4954)	0.5437 (0.4982)
<b>Highest degree</b>				
<i>Less than Bachelor's</i>	0.4395 (0.4964)	0.4459 (0.4971)	0.4408 (0.4966)	0.4320 (0.4954)
<i>Bachelor's</i>	0.3862 (0.4869)	0.3738 (0.4839)	0.3901 (0.4879)	0.3946 (0.4888)
<i>Graduate</i>	0.1743 (0.3794)	0.1803 (0.3845)	0.1691 (0.3749)	0.1734 (0.3786)
No. observations	8,498	2,817	2,797	2,884

Note. Standard deviations in parentheses.

Table 4.9.4 Main Results in Study 2.

	Outcome: Perc. Flatness		Attraction		
	Model 1	Model 2	Model 3	Model 4	Model 5
Estimation:	(OLS)	(OLS)	(OLS)	(2SLS)	(2SLS)
<i>No Information</i>	1.1993 (0.0286) [0.0000]	0.1494 (0.0250) [0.0000]	0.1968 (0.0399) [0.0000]		
<i>Flatter</i>	2.5003 (0.0244) [0.0000]	0.2748 (0.0254) [0.0000]	0.3490 (0.0406) [0.0000]		
<i>Perc. Flatness</i>				0.1097 (0.0101) [0.0000]	0.1384 (0.0160) [0.0000]
<i>Woman</i>			0.3001 (0.0382) [0.0000]		0.3795 (0.0686) [0.0000]
<i>No Information × Woman</i>			-0.0813 (0.0509) [0.1102]		
<i>Flatter × Woman</i>			-0.1227 (0.0517) [0.0178]		
<i>Perc. Flatness × Woman</i>					-0.0475 (0.0205) [0.0207]
No. observations	8,498	8,498	8,498	8,498	8,498
R <sup>2</sup>	0.511	0.032	0.048	0.031	0.045
KP Wald F statistic	-	-	-	5,273.5	1,622.6

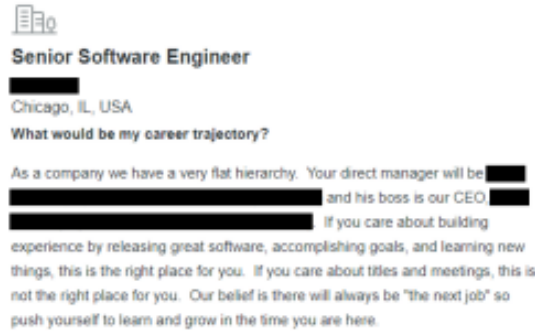
Note. Robust standard errors in parentheses and p-values in brackets.




## 4.10 Figures

Table 4.10.1 Featuring a flatter hierarchical structure in recruiting efforts.

(a) An example of a job posting that highlights a flatter organizational structure with fewer management layers. In this example, the names of the company, the CEO, and the direct manager are redacted.



(b) The growing trend of companies highlighting their flatter hierarchies in their job postings. In this figure, the observations are 6,424,963 firms, which posted a total of 156,705,625 jobs in the U.S. between 2010 and 2019. This sample was collected from Burning Glass Technologies, which provides a representative dataset of U.S. job postings (Cammeraat and Squicciarini 2021, Lee and Kim 2022).<sup>a</sup>



<sup>a</sup>To assess whether a firm's job posting features its hierarchical structure, we used the following list of phrases: "chain of command," "company hierarch\*," "corporate hierarch\*," "firm hierarch\*," "flat\* hierarch\*," "flat\* manag\*," "flat\* organi\*," "flat\* structur\*," "flatland," "formal hierarch\*," "hierarchical organi\*," "hierarchical structur\*," "holacracy," "layer\* of manag\*," "level\* of manag\*," "manag\* layer\*," "manag\* level\*," "no bosses," "no managers," "no middle manag\*," "non-hierarchical," "org\* hierarch\*," "org\* structur\*," "self-manag\*," "self-organi\*," "tall\* organi\*," and "tall\* structur\*," where the asterisk (\*) represents a meta-character for one or more instances of any alphabet letters. Among these phrases, those containing the word "flat" (e.g., "flat\* hierarch\*") or their equivalent terms (i.e., "holacracy," "no bosses," "no middle manag\*," "non-hierarchical," "self-manag\*," "self-organi\*"; Foss and Klein 2022, Lee and Edmondson 2017, Lee 2022, Reitzig 2022) were used to identify job postings specifically mentioning a flatter hierarchy. As this "flatter hierarchy" category does not include terms with "level" or "layer" (i.e., "layer\* of manag\*," "level\* of manag\*," "manag\* layer\*," and "manag\* level\*"), it does not capture job postings mentioning that the company, for example, has "very few layers of management." Thus, this category is likely to represent a conservative estimate for the proportion of firms highlighting their flatter hierarchies.

Table 4.10.2 Bar graphs with 95% confidence intervals for each dependent variable in Study 1.

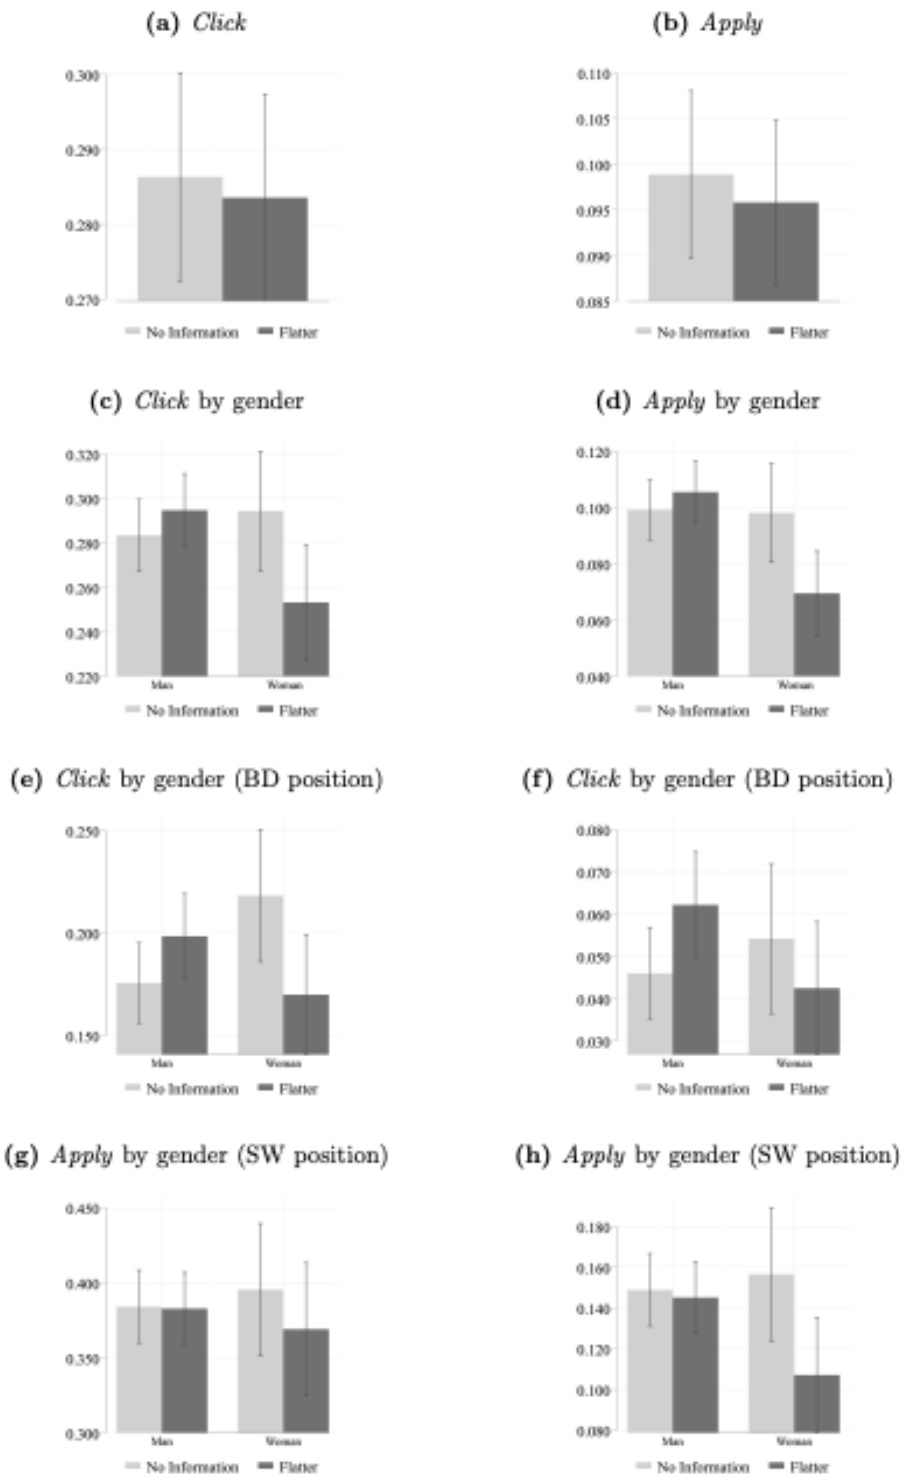
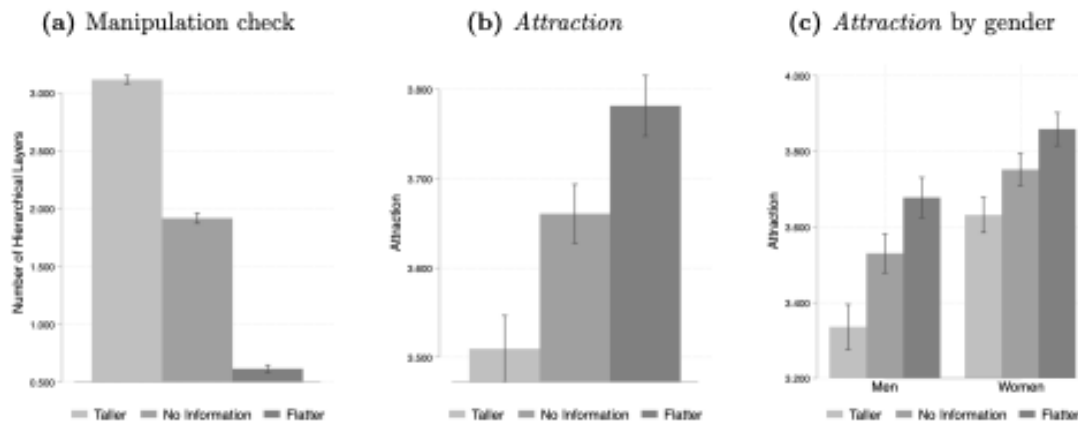


Table 4.10.3 Main bar graphs with 95% confidence intervals for Study 2.



## 4.11 Appendices

Table 4.11.1 Correlation matrix in Study 1.

(a) Full sample												
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
(1) <i>Click</i>	1.000											
(2) <i>Apply</i>	0.498	1.000										
(3) <i>Woman</i>	-0.015	-0.028	1.000									
(4) <i>Flatter</i>	-0.003	-0.005	-0.003	1.000								
(5) <i>Flatter × Woman</i>	-0.028	-0.037	0.649	0.393	1.000							
(6) <i>Software</i>	0.215	0.154	-0.085	0.001	-0.059	1.000						
(7) <i>Work Experience</i>	-0.056	-0.086	-0.058	0.003	-0.031	-0.309	1.000					
(8) <i>Manager</i>	-0.086	-0.081	-0.026	-0.002	-0.022	-0.451	0.348	1.000				
(9) <i>Willing to Relocate</i>	0.032	0.052	-0.060	-0.005	-0.040	0.083	-0.102	-0.018	1.000			
(10) <i>Bachelor's</i>	0.009	0.012	-0.024	0.018	-0.002	-0.061	0.023	0.054	0.006	1.000		
(11) <i>Masters</i>	0.027	0.022	0.020	-0.003	0.020	0.086	-0.007	-0.021	0.117	-0.078	1.000	
(12) <i>Doctorate</i>	0.010	0.010	-0.010	-0.001	-0.002	-0.011	0.072	0.008	-0.008	0.008	-0.099	1.000

(b) Business development (BD)											
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1) <i>Click</i>	1.000										
(2) <i>Apply</i>	0.453	1.000									
(3) <i>Woman</i>	0.008	-0.012	1.000								
(4) <i>Flatter</i>	0.001	0.017	0.004	1.000							
(5) <i>Flatter × Woman</i>	-0.021	-0.019	0.642	0.426	1.000						
(6) <i>Work Experience</i>	0.056	0.027	-0.136	0.004	-0.090	1.000					
(7) <i>Manager</i>	0.039	0.017	-0.088	-0.000	-0.069	0.243	1.000				
(8) <i>Willing to Relocate</i>	0.041	0.053	-0.079	-0.003	-0.048	-0.032	0.033	1.000			
(9) <i>Bachelor's</i>	0.002	-0.004	-0.007	0.026	0.014	0.016	0.032	0.025	1.000		
(10) <i>Masters</i>	0.022	0.024	-0.025	0.010	-0.004	0.032	0.042	0.118	-0.010	1.000	
(11) <i>Doctorate</i>	0.016	0.024	0.006	0.002	0.016	0.074	-0.005	-0.009	0.000	-0.094	1.000

(c) Software (SW)											
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1) <i>Click</i>	1.000										
(2) <i>Apply</i>	0.502	1.000									
(3) <i>Woman</i>	-0.001	-0.018	1.000								
(4) <i>Flatter</i>	-0.007	-0.020	-0.010	1.000							
(5) <i>Flatter × Woman</i>	-0.010	-0.037	0.653	0.357	1.000						
(6) <i>Work Experience</i>	-0.049	-0.123	-0.022	0.008	0.011	1.000					
(7) <i>Manager</i>	-0.021	-0.046	-0.047	-0.006	-0.031	0.259	1.000				
(8) <i>Willing to Relocate</i>	-0.007	0.033	-0.026	-0.007	-0.022	-0.151	0.005	1.000			
(9) <i>Bachelor's</i>	0.043	0.041	-0.055	0.008	-0.031	-0.023	0.026	-0.005	1.000		
(10) <i>Masters</i>	-0.001	-0.000	0.083	-0.017	0.058	0.003	-0.013	0.103	-0.138	1.000	
(11) <i>Doctorate</i>	0.010	0.005	-0.032	-0.004	-0.027	0.074	0.021	-0.005	0.016	-0.103	1.000

Table 4.11.2 Results using logistic regression in Study 1.

(a) Full sample

Outcome:	Click				Apply			
	Model 1 (All)	Model 2 (All)	Model 3 (Women only)	Model 4 (Men only)	Model 5 (All)	Model 6 (All)	Model 7 (Women only)	Model 8 (Men only)
<i>Flatter</i>	-0.0151 (0.0502) [0.7638]	0.0526 (0.0584) [0.3676]	-0.2055 (0.0981) [0.0362]	0.0526 (0.0585) [0.3681]	-0.0365 (0.0756) [0.6291]	0.0650 (0.0865) [0.4523]	-0.3725 (0.1581) [0.0184]	0.0650 (0.0865) [0.4523]
<i>Woman</i>		0.1444 (0.0800) [0.0710]				0.0821 (0.1198) [0.4934]		
<i>Flatter × Woman</i>		-0.2584 (0.1145) [0.0240]				-0.4375 (0.1801) [0.0151]		
Position FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
No. observations	8,167	8,167	2,195	5,972	8,167	8,167	2,195	5,972
Pseudo R <sup>2</sup>	0.039	0.040	0.039	0.040	0.038	0.040	0.043	0.037

Note. Robust standard errors in parentheses and p-values in brackets.

(b) Business development (BD)

Outcome:	Click				Apply			
	Model 9 (BD)	Model 10 (BD)	Model 11 (BD Women only)	Model 12 (BD Men only)	Model 13 (BD)	Model 14 (BD)	Model 15 (BD Women only)	Model 16 (BD Men only)
<i>Flatter</i>	0.0071 (0.0795) [0.9290]	0.1504 (0.0960) [0.1170]	-0.3105 (0.1432) [0.0301]	0.1504 (0.0960) [0.1170]	0.1545 (0.1402) [0.2702]	0.3189 (0.1669) [0.0561]	-0.2555 (0.2643) [0.3337]	0.3189 (0.1669) [0.0561]
<i>Woman</i>		0.2708 (0.1190) [0.0228]				0.1724 (0.2168) [0.4264]		
<i>Flatter × Woman</i>		-0.4609 (0.1723) [0.0075]				-0.5743 (0.3125) [0.0061]		
No. observations	4,132	4,132	1,264	2,868	4,132	4,132	1,264	2,868
Pseudo R <sup>2</sup>	0.000	0.002	0.004	0.001	0.001	0.003	0.002	0.003

Note. Robust standard errors in parentheses and p-values in brackets.

(c) Software (SW)

Outcome:	Click				Apply			
	Model 9 (SW)	Model 10 (SW)	Model 11 (SW Women only)	Model 12 (SW Men only)	Model 13 (SW)	Model 14 (SW)	Model 15 (SW Women only)	Model 16 (SW Men only)
<i>Flatter</i>	-0.0298 (0.0648) [0.6453]	-0.0054 (0.0738) [0.9421]	-0.1116 (0.1351) [0.4084]	-0.0054 (0.0738) [0.9421]	-0.1153 (0.0899) [0.1996]	-0.0299 (0.1014) [0.7680]	-0.4370 (0.1973) [0.0267]	-0.0299 (0.1014) [0.7680]
<i>Woman</i>		0.0477 (0.1077) [0.6581]				0.0592 (0.1454) [0.6842]		
<i>Flatter × Woman</i>		-0.1063 (0.1539) [0.4898]				-0.4071 (0.2217) [0.0664]		
No. observations	4,035	4,035	931	3,104	4,035	4,035	931	3,104
Pseudo R <sup>2</sup>	0.000	0.000	0.001	0.000	0.000	0.002	0.007	0.000

Note. Robust standard errors in parentheses and p-values in brackets.

Table 4.11.3 Results for various thresholds for Gender API's gender prediction in Study 1.

(a) Using the OLS regression

Outcome:	Click		Apply		Click		Apply	
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 5	Model 6
Threshold:	(70%)		(80%)		(90%)			
<i>Flatter</i>	0.0093 (0.0116) [0.4211]	0.0042 (0.0079) [0.5899]	0.0083 (0.0117) [0.4789]	0.0039 (0.0079) [0.6243]	0.0075 (0.0118) [0.5248]	0.0037 (0.0080) [0.6426]		
<i>Woman</i>	0.0242 (0.0158) [0.1257]	0.0022 (0.0103) [0.8337]	0.0224 (0.0159) [0.1600]	0.0035 (0.0105) [0.7408]	0.0203 (0.0161) [0.2070]	0.0049 (0.0107) [0.6453]		
<i>Flatter × Woman</i>	-0.0460 (0.0220) [0.0362]	-0.0287 (0.0140) [0.0408]	-0.0459 (0.0221) [0.0380]	-0.0304 (0.0141) [0.0315]	-0.0420 (0.0224) [0.0608]	-0.0308 (0.0144) [0.0323]		
No. observations	8,014	8,014	7,888	7,888	7,678	7,678		

Note. Robust standard errors in parentheses and p-values in brackets.

(b) Using the logistic regression

Outcome:	Click		Apply		Click		Apply	
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 5	Model 6
Threshold:	(70%)		(80%)		(90%)			
<i>Flatter</i>	0.0476 (0.0591) [0.4210]	0.0469 (0.0873) [0.5914]	0.0422 (0.0595) [0.4788]	0.0430 (0.0880) [0.6250]	0.0384 (0.0604) [0.5247]	0.0414 (0.0892) [0.6428]		
<i>Woman</i>	0.1253 (0.0807) [0.1206]	0.0275 (0.1220) [0.8217]	0.1163 (0.0816) [0.1542]	0.0431 (0.1227) [0.7253]	0.1054 (0.0824) [0.2010]	0.0596 (0.1232) [0.6288]		
<i>Flatter × Woman</i>	-0.2424 (0.1154) [0.0357]	-0.3839 (0.1824) [0.0353]	-0.2428 (0.1166) [0.0373]	-0.4085 (0.1842) [0.0265]	-0.2220 (0.1180) [0.0598]	-0.4090 (0.1855) [0.0274]		
No. observations	8,014	8,014	7,888	7,888	7,678	7,678		

Note. Robust standard errors in parentheses and p-values in brackets.

Table 4.11.4 Heterogenous effects of flatter hierarchy and gender in Study 1.

Outcome:	Click					Apply				
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8	Model 9	Model 10
<i>Flatter × Woman × Work Experience</i>	0.0009 (0.0030) [0.7565]					0.0006 (0.0018) [0.7335]				
<i>Flatter × Woman × Manager</i>		-0.0418 (0.0456) [0.3590]					-0.0097 (0.0267) [0.7162]			
<i>Flatter × Woman × Prior Employer's Hierarchy</i>				0.0143 (0.0104) [0.1691]				0.0073 (0.0062) [0.2392]		
<i>Flatter × Woman × Graduate Degree</i>									-0.0169 (0.0301) [0.5741]	
<i>Flatter × Woman × Elite Institution</i>										0.1160 (0.0816) [0.1550]
No. observations	6,894	8,137	4,114	8,167	8,167	6,894	8,137	4,114	8,167	8,167

Note. Robust standard errors in parentheses and p-values in brackets.

Table 4.11.5 Demographic attributes of MTurk workers in Study 2.

	<b>Mean</b>
<i>Woman</i>	0.584
<i>Manager</i>	0.247
<i>Work Experience</i>	
None	0.060
1-3 years	0.134
4-6 years	0.142
7-9 years	0.110
10+ years	0.554
<i>Education</i>	
Bachelor's	0.386
Graduate	0.174
Others	0.439
<i>Age</i>	37.665
<i>Race/Ethnicity</i>	
White	0.765
Black	0.102
Hispanic	0.072
Asian	0.086
Others	0.023
<i>Employment</i>	
Full-time	0.592
Part-time	0.173
Unemployed	0.235
<i>Income</i>	
Less than \$40K	0.251
\$40-60K	0.215
\$60-80K	0.146
\$80-100K	0.175
More than \$100	0.213
<i>Industry</i>	
Healthcare	0.130
IT	0.101
Education	0.120
Retail	0.082
Others	0.567
No. observations	8,498



Table 4.11.6 Measurement of primary variables in Study 2.

Variable	Measurement
<b>Dependent</b>	
<i>Attraction</i>	Average score of the five-point Likert scales for the following five items: <ul style="list-style-type: none"> <li>· “For me, Exogeny would be a good place to work.”</li> <li>· “I would not be interested in Exogeny except as a last resort.” (reverse-coded)</li> <li>· “Exogeny is attractive to me as a place of employment.”</li> <li>· “I am interested in learning more about Exogeny.”</li> <li>· “A job at Exogeny is very appealing to me.”</li> </ul>
<b>Independent</b>	
<i>No Information</i>	1 if no mention of the organization’s hierarchical structure; 0 otherwise.
<i>Flatter</i>	1 if flatter condition; 0 otherwise.
<b>Individual attributes</b>	
<i>Woman</i>	1 if the subject self-identified as female; 0 otherwise.
<b>Boundary condition</b>	
<i>Woman Founder</i>	1 if the subject received the female founder/CEO condition; 0 otherwise.

Table 4.11.7 Measurement of mechanisms in Study 2.

<b>Mechanism</b>	<b>Measurement</b>
<i>Opportunity</i>	Average score of the five-point Likert scales for the following two items: · “It would be possible for someone like me to receive a pay raise at Exogeny.” · “It would be possible for someone like me to secure a promotion at Exogeny.”
<i>Autonomy</i>	Average score of the five-point Likert scales for the following four items: · “Exogeny would support me to work autonomously instead of requiring me to rely on managers to guide my work.” · “Exogeny’s leaders believe that the best results occur when workers decide on their own how to spend their time at work.” · “At Exogeny, individuals make decisions on their own without the need for approval from their supervisors.” · “At Exogeny, employee input (rather than CEO and top management decisions) plays a major role in determining what actions the firm takes.”
<i>Informality</i>	Average score of the reverse-coded five-point Likert scales for the following four items: · “Exogeny has clearly defined procedures for evaluating performance.” · “Exogeny has clearly defined procedures for determining pay.” · “Exogeny has clearly defined procedures for determining promotion.” · “Exogeny has clearly defined procedures for resolving conflicts among employees.”
<i>Fairness</i>	Average score of the five-point Likert scales for the following four items: · “Exogeny has fair procedures for evaluating performance.” · “Exogeny has fair procedures for determining pay.” · “Exogeny has fair procedures for determining promotion.” · “Exogeny has fair procedures for resolving conflicts among employees.”
<i>Fit</i>	The score of the five-point Likert scales on: · “I would fit in well at Exogeny.”
<i>Workload</i>	The score of the five-point Likert scales on: · “My workload at Exogeny would be heavy.”
<i>Competition</i>	The score of the five-point Likert scales on: · “Exogeny has a very competitive work environment.”
<i>Firm Success</i>	The score of the five-point Likert scales on: · “Exogeny is likely to succeed.”

Table 4.11.8 Correlation matrix in Study 2.

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)
(1) Attraction	1.000																
(2) Flatter	0.068	1.000															
(3) No Information	0.007	-0.502	1.000														
(4) Teller	-0.106	-0.505	-0.493	1.000													
(5) Woman	0.120	-0.012	0.008	0.005	1.000												
(6) Flatter x Woman	0.107	0.687	-0.345	-0.347	0.416	1.000											
(7) No Information x Woman	0.052	-0.352	0.701	-0.346	0.414	-0.242	1.000										
(8) Teller x Woman	-0.010	-0.352	-0.344	0.698	0.415	-0.242	-0.241	1.000									
(9) Manager	-0.018	-0.019	0.001	0.018	-0.119	-0.056	-0.045	-0.047	1.000								
(10) No Experience	-0.034	-0.007	0.000	0.007	-0.003	-0.010	0.001	0.006	-0.130	1.000							
(11) 1-3 years	-0.050	0.028	-0.022	-0.006	0.019	0.023	-0.013	0.003	-0.135	-0.099	1.000						
(12) 4-6 years	-0.026	0.004	-0.001	-0.003	0.021	0.031	-0.015	0.000	-0.029	-0.102	-0.160	1.000					
(13) 7-9 years	-0.013	-0.005	-0.007	0.013	-0.002	-0.010	0.013	0.004	-0.001	-0.088	-0.138	-0.143	1.000				
(14) 10+ years	0.078	-0.015	0.020	-0.004	-0.025	-0.027	0.014	-0.007	0.176	-0.281	-0.439	-0.453	-0.392	1.000			
(15) Bachelor's	-0.050	0.012	0.006	-0.018	-0.036	-0.006	-0.012	-0.028	0.068	-0.061	0.018	0.022	0.050	-0.030	1.000		
(16) Graduate	-0.055	-0.002	-0.010	0.011	0.004	0.008	-0.015	0.002	0.089	-0.083	-0.047	-0.011	-0.003	0.081	-0.364	1.000	
(17) Less than Bachelor's	0.091	-0.011	0.002	0.009	0.032	-0.000	0.015	0.026	-0.135	0.123	0.019	-0.013	-0.047	-0.033	-0.702	-0.407	1.000

Table 4.11.9 Post-hoc mechanism tests for Study 2.

(a) Using the OLS regression by condition

Outcome:	Opportunity	Autonomy	Informality	Fairness	Fit	Workload	Competition	Firm Success	
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8	
<i>No Information</i>	-0.0491 (0.0230) [0.0325]	0.4484 (0.0240) [0.0000]	0.0869 (0.0267) [0.0011]	0.0461 (0.0196) [0.0189]	0.1652 (0.0266) [0.0000]	0.0132 (0.0228) [0.5610]	0.0132 (0.0235) [0.0000]	-0.2780 (0.0235) [0.0000]	0.0325 (0.0233) [0.1635]
<i>Flatter</i>	-0.3153 (0.0243) [0.0000]	0.9517 (0.0238) [0.0000]	0.1182 (0.0265) [0.0000]	-0.0006 (0.0200) [0.9770]	0.2542 (0.0266) [0.0000]	0.0315 (0.0223) [0.1582]	-0.2871 (0.0252) [0.0000]	0.0720 (0.0234) [0.0021]	
No. observations	8,498	8,498	8,498	8,498	8,498	8,498	8,498	8,498	
R <sup>2</sup>	0.037	0.172	0.015	0.017	0.025	0.003	0.021	0.013	

Note. Robust standard errors in parentheses and p-values in brackets.

(b) Using the OLS regression by condition and gender

Outcome:	Opportunity	Autonomy	Informality	Fairness	Fit	Workload	Competition	Firm Success
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
<i>No Information</i>	0.0162 (0.0359) [0.6527]	0.5169 (0.0383) [0.0000]	0.0974 (0.0410) [0.0175]	0.0593 (0.0310) [0.0559]	0.2544 (0.0429) [0.0000]	-0.0693 (0.0360) [0.0542]	-0.2399 (0.0365) [0.0000]	0.0465 (0.0373) [0.2130]
<i>Flatter</i>	-0.2130 (0.0369) [0.0000]	1.0535 (0.0379) [0.0000]	0.1128 (0.0402) [0.0050]	0.0148 (0.0312) [0.8357]	0.3510 (0.0424) [0.0000]	-0.0661 (0.0347) [0.0570]	-0.1973 (0.0384) [0.0000]	0.1108 (0.0368) [0.0026]
<i>Woman</i>	0.0497 (0.0338) [0.1410]	0.1494 (0.0388) [0.0001]	0.0545 (0.0378) [0.1496]	0.0139 (0.0288) [0.6285]	0.2193 (0.0405) [0.0000]	-0.2012 (0.0321) [0.0000]	0.0189 (0.0340) [0.5780]	0.2496 (0.0346) [0.0000]
<i>No Information × Woman</i>	-0.1109 (0.0467) [0.0176]	-0.1167 (0.0491) [0.0176]	-0.0179 (0.0540) [0.7396]	-0.0225 (0.0400) [0.5735]	-0.1520 (0.0544) [0.0033]	0.1406 (0.0463) [0.0024]	-0.0646 (0.0477) [0.1755]	-0.0246 (0.0474) [0.6041]
<i>Flatter × Woman</i>	-0.1767 (0.0490) [0.0003]	-0.1738 (0.0487) [0.0004]	0.0105 (0.0534) [0.8447]	-0.0264 (0.0407) [0.5168]	-0.1637 (0.0544) [0.0026]	0.1654 (0.0453) [0.0003]	-0.1557 (0.0509) [0.0022]	-0.0623 (0.0474) [0.1885]
No. observations	8,498	8,498	8,498	8,498	8,498	8,498	8,498	8,498
R <sup>2</sup>	0.039	0.174	0.016	0.017	0.030	0.008	0.023	0.029

Note. Robust standard errors in parentheses and p-values in brackets.

(c) Using 2SLS approach by condition

Outcome:	Opportunity	Autonomy	Informality	Fairness	Fit	Workload	Competition	Firm Success
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
<i>Perc. Flatness</i>	-0.1275 (0.0097) [0.0000]	0.3807 (0.0097) [0.0000]	0.0468 (0.0106) [0.0000]	-0.0009 (0.0080) [0.9137]	0.1011 (0.0106) [0.0000]	0.0126 (0.0089) [0.1572]	-0.1129 (0.0101) [0.0000]	0.0288 (0.0093) [0.0020]
No. observations	8,498	8,498	8,498	8,498	8,498	8,498	8,498	8,498
R <sup>2</sup>	0.029	0.092	0.018	0.017	0.021	0.001	0.017	0.011

Note. Robust standard errors in parentheses and p-values in brackets.

(d) Using 2SLS approach by condition and gender

Outcome:	Opportunity	Autonomy	Informality	Fairness	Fit	Workload	Competition	Firm Success
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
<i>Perc. Flatness</i>	-0.0861 (0.0147) [0.0000]	0.4186 (0.0152) [0.0000]	0.0443 (0.0160) [0.0056]	0.0052 (0.0124) [0.6740]	0.1385 (0.0168) [0.0000]	-0.0259 (0.0138) [0.0612]	-0.0764 (0.0154) [0.0000]	0.0442 (0.0146) [0.0025]
<i>Woman</i>	0.1793 (0.0636) [0.0048]	0.2515 (0.0684) [0.0002]	0.0365 (0.0702) [0.6030]	0.0307 (0.0532) [0.5646]	0.3113 (0.0726) [0.0000]	-0.3034 (0.0596) [0.0000]	0.1429 (0.0650) [0.0279]	0.2972 (0.0628) [0.0000]
<i>Perc. Flatness × Woman</i>	-0.0717 (0.0196) [0.0003]	-0.0648 (0.0197) [0.0010]	0.0050 (0.0213) [0.8161]	-0.0105 (0.0163) [0.5189]	-0.0633 (0.0217) [0.0035]	0.0654 (0.0181) [0.0003]	-0.0633 (0.0204) [0.0019]	-0.0246 (0.0189) [0.1920]
No. observations	8,498	8,498	8,498	8,498	8,498	8,498	8,498	8,498
R <sup>2</sup>	0.030	0.094	0.019	0.017	0.024	0.002	0.019	0.026

Note. Robust standard errors in parentheses and p-values in brackets.

Table 4.11.10 Results for other potential mechanisms in Study 2.

	Outcome: <u>Attraction</u>	<u>Attraction</u>	<u>Attraction</u>
	Model 1	Model 2	Model 3
<i>No Information × Woman × Woman Founder/CEO</i>	0.0274 (0.1054) [0.7951]		
<i>Flatter × Woman × Woman Founder/CEO</i>	0.0243 (0.1074) [0.8210]		
<i>No Information × Woman × Share of Women</i>		-0.9281 (0.3225) [0.0040]	
<i>Flatter × Woman × Share of Women</i>		-0.3516 (0.3521) [0.3180]	
<i>No Information × Woman × Founder Competence</i>			0.0539 (0.0546) [0.3236]
<i>Flatter × Woman × Founder Competence</i>			0.0680 (0.0554) [0.2197]
No. observations	8,498	8,498	8,498

*Note.* Robust standard errors in parentheses and p-values in brackets.

Figure 4.11.1 Email invitation to job candidates in Study 1.

---

Dear *{Name of job candidate}*,

We found your profile on ZipRecruiter. Based on your credentials, we invite you to apply for a position in [business development/software engineering] here at *{Partner company's name}*.

*{ "Flatter" condition only: We think you'll be interested to know that we have a flat organizational structure. This means that we have fewer levels of management than similarly sized startups in our industry. }*

We are a telemedicine company specializing in online therapy. We founded *{Partner company's name}* in *{Founding year}*, and have since expanded to *{Number of employees}* full-time employees. We are looking to grow significantly.

To learn more about us and to apply for this position, please click this link to our job posting: *{Hyperlink to the company's application web page}*.

---

Figure 4.11.2 Recruiting material in Study 2 with no mention of the organization's hierarchical structure or the founder/CEO's gender.

---

**Part-Time, Remote Copyeditor**

- We are seeking a part-time, remote copyeditor to help edit our reports and presentations.

**About Us**

- In 2015, we founded Exogeny to provide data-driven HR analytics.

**Job Responsibilities**

- Proofread reports and presentations we prepare for external audiences
- Provide feedback on the clarity of writing style and graphics

**Requirements**

- Exceptional attention to detail
- English fluency
- Ability to meet deadlines

---

Figure 4.11.3 The interventions in Study 2 inserted in the "About Us" section for the organization's hierarchical structure or the founder/CEO's gender.







Manipulation	Condition	Description
Organizational structure	<i>Flatter</i>	"A defining feature is our flat organizational structure—we have fewer levels of management than similarly sized startups in our industry."
		<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;">  <p>Our structure</p> </div> <div style="text-align: center;">  <p>Our competitors' structure</p> </div> </div>
	<i>Taller</i>	"A defining feature is our tall organizational structure—we have more levels of management than similarly sized startups in our industry."
		<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;">  <p>Our structure</p> </div> <div style="text-align: center;">  <p>Our competitors' structure</p> </div> </div>
Founder/CEO's gender	<i>Woman</i>	<div style="text-align: center;">  <p>Jessica Chandler</p> </div>
	<i>Man</i>	<div style="text-align: center;">  <p>Michael Chandler</p> </div>

Figure 4.11.4 Bar Graphs with 95% confidence intervals for mechanisms in Study 2.

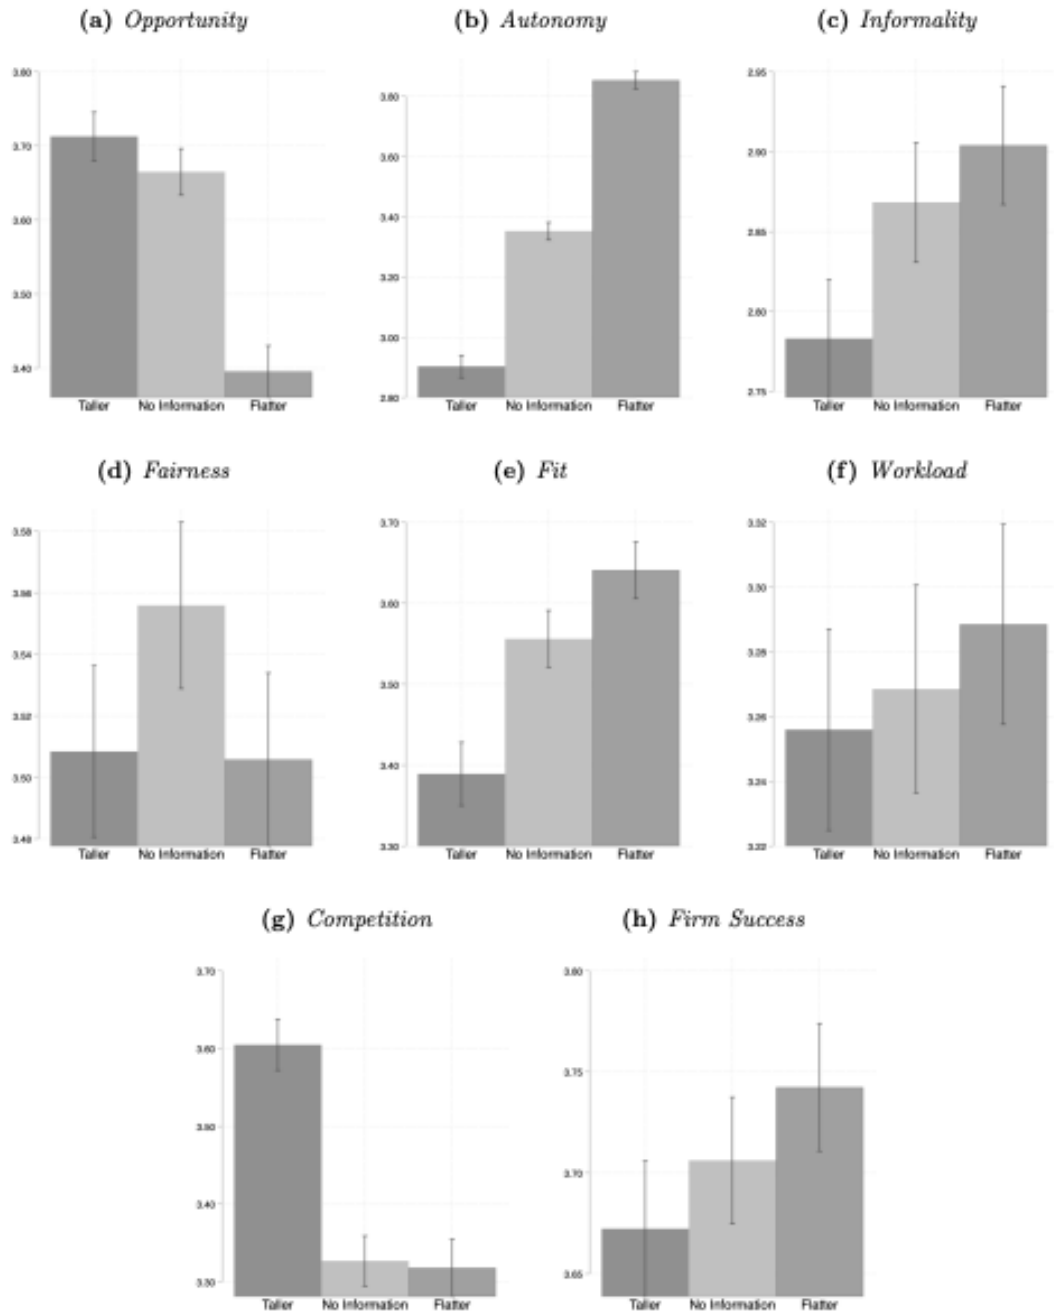
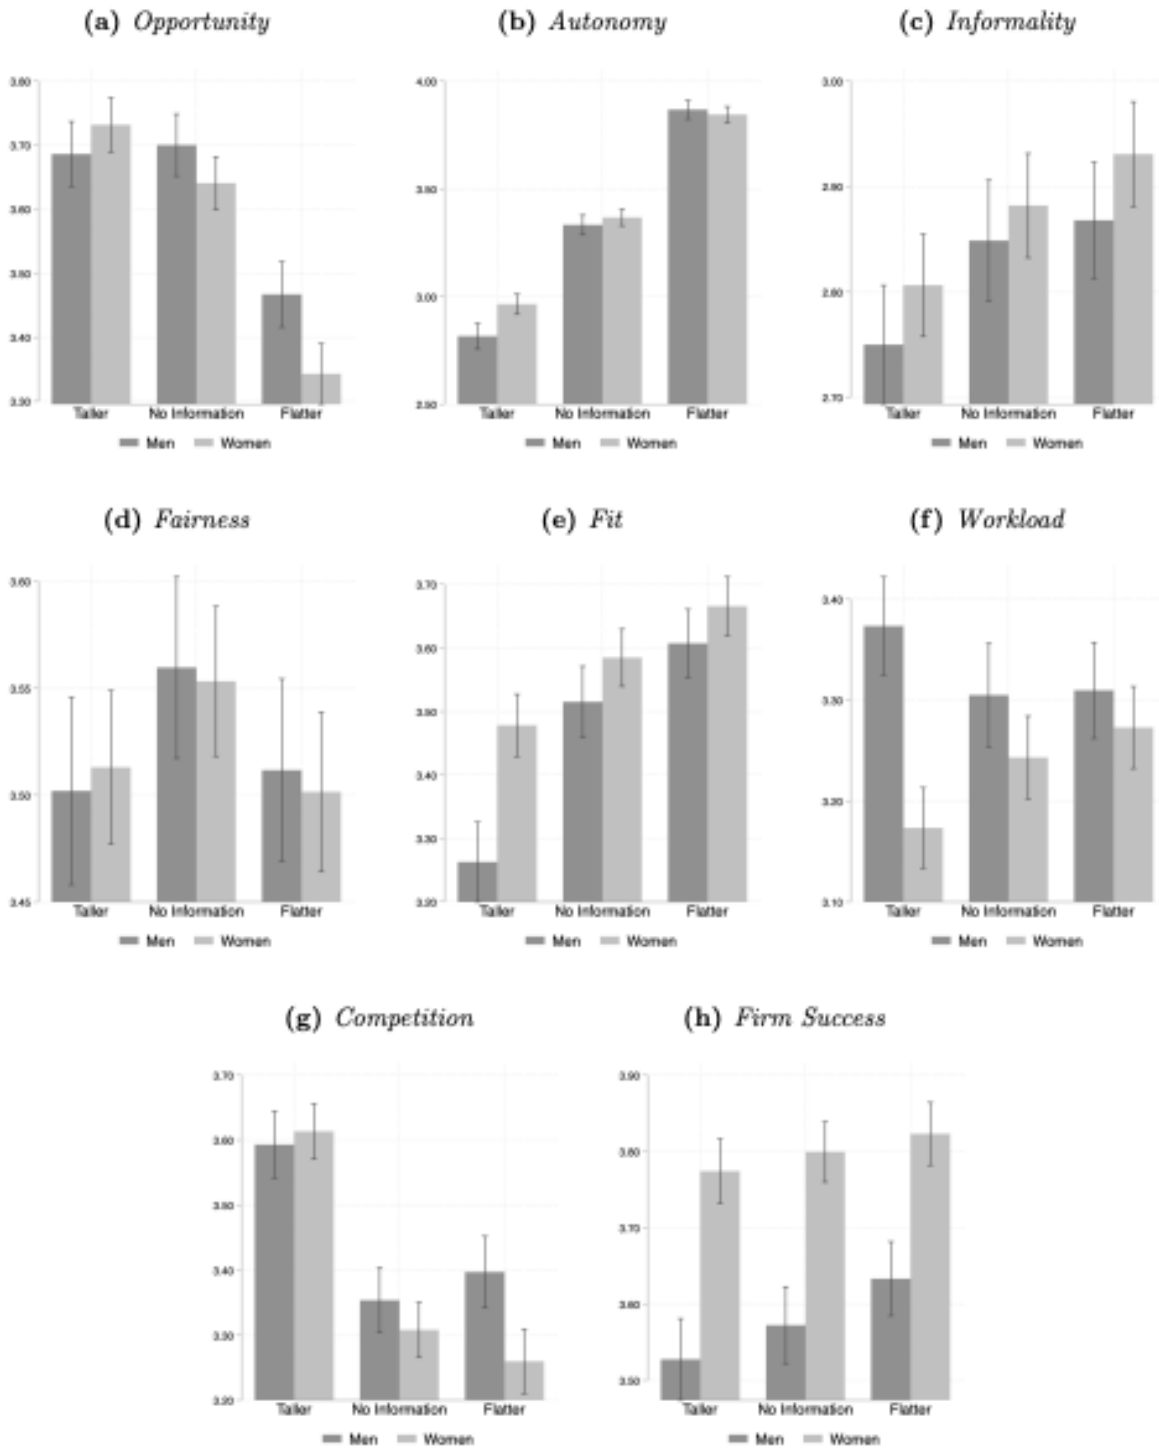




Figure 4.11.5 Bar graphs with 95% confidence intervals for mechanisms by gender in Study 2.



#### 4.11.1 Heterogeneous effects of flatter hierarchy and gender in Study 1

Here, we explore potential heterogeneous effects by respondents' individual characteristics in Study 1. These results are summarized in Table A4. First, as more experienced workers may have a better understanding of how taller and flatter hierarchies function, we examine whether job seekers' years of work experience (measured by the number of years since their first post-undergraduate job began in their ZipRecruiter profile) affect the strength of their attraction to firms with flatter hierarchies. However, in Models 1 (for Click) and 6 (for Apply), we find no evidence that the effects are moderated by the job seeker's years of work experience.

Second, as managers and entry-level employees may have different perceptions of formal hierarchy given their current hierarchical ranks, we measured whether job seekers' most recent job title in their ZipRecruiter profile includes a keyword indicating a managerial position (i.e., "president," "vice president," "chief . . . officer," "manager," "director," "head," "lead," and their abbreviations, such as "VP," "CEO," or "mgr"). But, Models 2 (for Click) and 7 (for Apply) do not show results to support this argument.

Third, as the most recent employer's hierarchy may influence how job seekers view flatter or taller hierarchies, we counted the number of hierarchical levels in their most recent employer. To compute this measure, we collected additional data from LinkedIn and matched this data with our ZipRecruiter dataset based on the job seeker's full name and his/her most recent employer's name. For each of these most recent employers, we gathered all employees with publicly available profiles on LinkedIn as of 2022 and drew upon the methodology in prior studies (e.g., Lee 2022, Lee and Csaszar 2020) to categorize each employee's job title into a hierarchical level. We then counted the number of distinct hierarchical levels and included this measure of their prior

employer's hierarchy in a three-way interaction with Woman and Flatter. In Models 3 (for Click) and 8 (for Apply), we do not find evidence to support this three-way interaction.

Finally, we assessed whether job seekers' educational background, measured by whether they have a graduate degree or received a bachelor's degree from an elite institution (i.e., U.S. News Top 50 universities in 2022), moderates their attraction to the flatter employers. Models 4 (for Click) and 9 (for Apply) offer no evidence that having a graduate degree moderates the relationship. However, Models 5 (for Click) and 10 (for Apply) suggest that women with a bachelor's degree from an elite institution are somewhat more likely to prefer flatter firms ( $p = .0135$  and  $p = .1550$ , respectively). One possible explanation for this result is that women with elite undergraduate degrees believe their elite credentials will shield them from marginalization or discrimination in flatter organizations. Because these post-hoc analyses do not allow us to directly test our proposed mechanisms (e.g., perceptions of autonomy or career progression), we conduct a survey experiment.

#### 4.11.2 Examinations of Mechanisms Study 1

Here, we discuss in detail the regression results for the post-hoc mechanism analyses. These results are reported in Table A9 and shown visually using bar charts in Appendices A13 and A14. We also present the results from the ANOVA analyses in parentheses.

**1. Opportunity:** Model 1 shows that subjects, on average, perceive less opportunity for career advancement in flatter employers than in taller ones (see panel (a) in Figure A13,  $F(2, 8, 495) = 103.44$ ,  $p = .000$ ). Model 9 demonstrates that, relative to men, women expect this opportunity deficit in the flatter hierarchies to be even more severe (see panel (a) in Figure A14,  $F(2, 8, 492) = 6.29$ ,  $p = .002$ ).

**2. Autonomy:** Model 2 indicates that subjects, on average, anticipate flatter employers to offer more autonomy than taller ones (see panel (b) in Figure A13,  $F(2, 8, 495) = 863.57, p = .000$ ). Model 10 shows that the effect is moderated by gender: that is, compared to men, women perceive the flatter employer to offer them less autonomy (see panel (b) in Figure A14,  $F(2, 8, 492) = 7.09, p = .001$ ).

**3. Informality:** Model 3 suggests that, on average, subjects perceive flatter employers to be relatively more informal than taller ones (see panel (c) in Figure A13,  $F(2, 8, 495) = 10.78, p = .000$ ). However, Model 11 implies that gender does not significantly moderate this effect (see panel (c) in Figure A14,  $F(2, 8, 492) = 0.16, p = .850$ ).

**4. Fairness:** Model 4 shows that flatter and taller employers are perceived as less fair than the No Information condition (see panel (d) in Figure A13,  $F(2, 8, 495) = 3.96, p = .019$ ), but the difference between the flatter and taller employers is not statistically different. Furthermore, Model 12 demonstrates that the effects are not significantly moderated by gender (see panel (d) in Figure A14,  $F(2, 8, 492) = 0.16, p = .855$ ).

**5. Fit:** Model 5 indicates that subjects, on average, perceive they will fit into flatter organizations better than into taller ones (see panel (e) in Figure A13,  $F(2, 8, 495) = 47.5, p = .000$ ). In addition, Model 13 reveals that the effects are significantly moderated by gender: that is, compared to men, women perceive a relatively less fit in flatter firms (see panel (e) in Figure A14,  $F(2, 8, 492) = 5.38, p = .005$ ).

**6. Workload:** Model 6 demonstrates that, on average, formal hierarchy does not significantly affect how subjects perceive their workload (see panel (f) in Figure A13,  $F(2, 8, 495) = 1.06, p = .340$ ). Interestingly, Model 14 shows that women perceive that they will have (absolutely) less workload in taller firms, while men perceived they would have (absolutely) less workload in flatter

firms. ANOVA analysis suggests the interaction between hierarchy and gender on perceived workload is significant (see panel (f) in Figure A14,  $F(2, 8, 492) = 7.38, p = .001$ ).

**7. Competition:** As shown in Model 7, flatter employers are perceived to have less competitive work environments than taller ones (see panel (g) in Figure A13,  $F(2, 8, 495) = 87.5, p = .000$ ). Furthermore, Model 15 reveals that, relative to men, women perceive even less competition in flatter employers. Again, ANOVA analysis finds significant interaction between gender and hierarchy on competition (see panel (g) in Figure A14,  $F(2, 8, 492) = 5.08, p = .006$ ).

**8. Firm Success:** Model 8 shows that subjects, on average, perceive the Flatter employer as significantly more likely to succeed than both the No Information and Taller ones. The difference between the No Information and Taller conditions is insignificant. A one-way ANOVA analysis also finds that hierarchy significantly affects the subjects' perceptions of firm success (see panel (h) in Figure A13,  $F(2, 8, 495) = 4.57, p = .010$ ). However, the effects are not significantly moderated by gender (see panel (h) in Figure A14,  $F(2, 8, 492) = 0.77, p = .464$ ).