
Increasing Workplace Diversity

Evidence from a Recruiting Experiment at a Fortune 500 Company

Jeffrey A. Flory
Andreas Leibbrandt
Christina Rott
Olga Stoddard

ABSTRACT

While many firms have set ambitious goals to increase diversity in their ranks, there is a dearth of empirical evidence on effective ways to reach them. We use a natural field experiment to test several hypotheses on effective means to attract minority candidates for top professional careers. By randomly varying the content in recruiting materials of a major financial services corporation with more than 10,000 employees, we find that signaling explicit interest in employee diversity more than doubles the interest in openings among racial minority candidates, as well as the likelihood that they apply and are selected. Impacts on gender diversity are less sharp and generally not significant.

Jeffrey A. Flory is Associate Professor of Economics at Claremont McKenna College. Andreas Leibbrandt is Professor of Economics at Monash University. Christina Rott is Assistant Professor of Economics at Vrije Universiteit Amsterdam. Olga Stoddard is Assistant Professor of Economics at Brigham Young University (olga.stoddard@byu.edu). The authors thank Rebecca Jack for excellent research assistance and gratefully acknowledge financial support from the Bias Interrupters Working Group. The authors also thank participants of SOCCAM, ESA AFE, IMEBESS, and Science of Diversity and Inclusion Initiative (SODI) conferences, as well as seminars at University of California (San Diego), Brigham Young University, Claremont McKenna College, University of Melbourne, Maastricht University, Amsterdam Business School, Vrije Universiteit of Amsterdam, and Tilburg University for their helpful comments on earlier drafts of the paper. The authors are listed in alphabetic order. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. Institutional Review Board (IRB) approval was obtained to conduct this research involving the collection of data on human subjects. The data are available at olgastoddard.byu.edu/research.

[Submitted May 2018; accepted April 2019]; doi:10.3368/jhr.56.1.0518-9489R1

JEL Classification: J01 and D01

ISSN 0022-166X E-ISSN 1548-8004 © 2021 by the Board of Regents of the University of Wisconsin System

 Supplementary materials are freely available online at: <http://uwpress.wisc.edu/journals/journals/jhr-supplementary.html>

I. Introduction

Understanding and eliminating barriers to career progress based on bias against people on the basis of their race, gender, and other social identity, rather than on skills and ability, is an important issue in the U.S. labor market. Removing social identity-based frictions in the flow of human capital to its most productive uses increases economic efficiency, making demographic diversity an increasing priority for firms. Estimates suggest companies spend nearly \$10 billion a year on diversity and inclusion training initiatives (Hansen 2003), with some dedicating as much as \$200 million a year or more to diversity programs (Catalyst 2005), and the Chief Diversity Officer has now become an established leadership position in the country's largest firms.¹ However, despite significant education gains among underrepresented groups, and substantial resources devoted to enhancing employee diversity in high-profile occupations, many firms still struggle to increase representation by individuals from minority groups. Through a field experiment in a major U.S. corporation, we show how simple changes in presentation of recruitment materials can have powerful impacts on selection into high-profile opportunities by people in underrepresented groups, and ethnic minorities in particular.

The managerial landscape in many U.S. companies continues to resemble monocultures. Almost 90 percent of Fortune 500 CEOs are white males, while less than 4 percent are African American or Hispanic, and less than 6 percent are women.² Among all U.S. companies with 100 or more employees, the proportion of Black men in management barely increased from 3 percent in 1985 to 3.3 percent in 2014, and the proportion of white women has stayed mostly flat since 2000 at under 30 percent (Dobbin and Kalev 2016). These percentages have remained remarkably low, despite significant advances in educational achievements of racial minorities and women over the last few decades. African Americans and women, for example, account for ever larger proportions of MBA-holders in the United States, rising from 4 percent in 1990 to 14 percent in 2015 for African Americans and from 22 percent in 1980 to 47 percent in 2014 for women (National Center for Education Statistics 2018). This disparity suggests important barriers in the career trajectories for racial minorities and women, but it also suggests significant opportunities for organizations to increase demographic diversity among employees. However, the key question remains how.

Nearly all Fortune 500 companies and almost half of all mid-size companies in the United States have programs to enhance employee diversity (Dobbin and Kalev 2016). These programs contain a range of company-specific initiatives to reduce discrimination and unintentional biases in employee hiring and promotion and to provide resources for underrepresented groups, including scholarships, mentoring,

1. Korn Ferry, the top executive recruiting firm, states that about 60 percent of the Fortune 500 companies have Chief Diversity Officers, or the equivalent: <https://www.kornferry.com/functions/div><https://www.kornferry.com/functions/diversity-officers> (accessed May 29, 2020).

2. See, for example, <http://fortune.com/2015/12/23/2015-women-fortune-500-ceos>, http://www2.ucsc.edu/whorulesamerica/power/rise_and_fall_of_diversity.html, <http://fortune.com/2016/06/06/lessons-fortune-500>, <http://fortune.com/2013/05/09/women-ceos-in-the-fortune-500> (accessed April 21, 2020).

and support groups. However, much remains unknown about the actual impacts of programs aiming to foster workplace diversity. This makes it very difficult to evaluate their cost-effectiveness. To make matters worse, there is suggestive evidence that some diversity programs can actually backfire (Gilbert and Ivancevich 2000; McKay and Avery 2005; Kalev, Dobbin, and Kelly 2006; Apfelbaum, Stephens, and Reagans 2016), increasing the urgency for rigorous causal analysis of these different approaches.

In order to better understand how minority candidates can be attracted for future management positions, we conduct a natural field experiment in one of the largest firms in the financial services industry—a setting widely perceived to lack both racial and gender diversity. By randomizing the content of recruiting information across individuals, we exogenously vary whether and how the employer signals that it values diversity among its employees. We also systematically vary whether the diversity statements are backed by facts in order to assess the importance of supporting information when using this type of approach to raise employee diversity. This is important since some organizations may be unable to furnish evidence on stated diversity values, and it is unclear whether individuals respond to unsupported diversity statements or instead view it as “cheap talk.” Finally, we compare the impact of statements that directly target candidates from underrepresented groups to statements that use a less direct appeal—since there may be settings in which an overt approach to diversity recruiting is impractical or undesirable, and since there is suggestive evidence that some types of direct signals can actually push away skilled employees from racial minority groups (Leibbrandt and List 2018).

This work is related to a growing literature on understanding and eliminating barriers to labor market entry and career progress based on race, gender, and other social identities (Hinton et al. 2010; Giuliano, Levine, and Leonard 2011; Bertrand and Duflo 2017; Blau and Kahn 2017). A rich vein of this literature includes field experiments and studies on gender differences in willingness to enter competitive environments (Buser, Neiderle, and Oosterbeek 2014; Flory, Leibbrandt, and List 2014; Preece and Stoddard 2015, 2016; Flory et al. 2018) and how affirmative action policies and quotas can reduce the gender gap (Balafoutas and Sutter 2012; Niederle, Segal, and Vesterlund 2013; Leibbrandt et al. 2018; Ibanez and Riener 2018). While there are several related studies showing that the way a job is described can affect the applicant pool (Ashraf, Bandiera, and Lee 2014; Dal Bó, Finan, and Rossi 2013; Marinescu and Wolthoff 2016), no prior research has investigated how signaling a firm’s valuation of diversity affects interest and applications by people from underrepresented ethnic/racial minority groups.

The paper arguably most closely related to ours is by Del Carpio and Guadalupe (2018). In a series of two field experiments in Lima and Mexico City, the authors investigate whether varying information provided to potential applicants to a coding program is effective in attracting women. Randomized messages contained information about role models, career prospects, and peer network. The authors find that exposure to the message significantly increased application rates by women, with the statement about role models being the most effective. While this paper sheds light on effective ways to narrow the gender gap in historically male-dominated fields like tech, it does not explore racial or ethnic diversity, which is at the core of our study.

II. Hypotheses and Experimental Design

A. Research Hypotheses

Underrepresented groups may be dissuaded by workplaces they perceive as lacking in demographic diversity, such as race and gender, for several reasons. From a pecuniary perspective, avoiding certain work environments may be a rational payoff-maximizing response to income risks from working in settings with little or no representation of one's own demographic group. This may be due to expectations of discrimination by supervisors or unconscious biases in performance evaluations, causing an individual to believe her effort will be less rewarded in these settings. It may also be a strategic avoidance of risk that actual performance might become lower in certain work settings. Working in an environment where there are few or no other individuals of the same identity group is likely to make that identity salient, and several studies have shown that emphasizing social identities such as race and gender can undermine performance, lower effort and expected success, and increase anxiety—often referred to as “stereotype threat” (Steele and Aronson 1995; see Bertrand and Duflo 2017 for a review). This may cause some from less represented groups to seek work settings where their social identity is less salient. From a nonpecuniary perspective, individuals may have a preference for work settings that include others from their own group, or workers from underrepresented groups may prefer environments with greater representation from a variety of different demographics.

This suggests that being perceived as lacking in diversity of social identities may cause an employer to have difficulty attracting individuals from underrepresented groups. However, it also suggests a firm may be able to increase its ability to attract such candidates by presenting itself as diverse or as placing high value on the diversity of its workforce. This leads to our first main hypothesis.

Hypothesis 1: Projecting diversity among employees as an organizational value and priority will increase interest among underrepresented groups and achieve greater demographic diversity in recruiting outcomes.

However, in some circles there is concern about an approach to workplace diversity that is too overt. There have also been reports of instances of backlash against diversity efforts, and discontent among individuals excluded from these efforts—typically, ethnic-majority males.³ Some firms and organizations may therefore be reluctant to appear as if they are intentionally trying to court underrepresented groups, while still wanting to increase their rates of recruitment.

One approach is to use an appeal to cognitive diversity, by encouraging interest among individuals from a variety of *educational* backgrounds. The conceptual basis for this type of approach is threefold. First, appealing to fields of study other than those typically associated with the industry, occupation, or firm can have a mechanical effect if those fields have greater proportions of underrepresented groups. Second, expressing an interest in diversity of educational background may be interpreted by job-seekers as a signal of an openness or desire for more general diversity among employees (including dimensions of social identity), which may attract candidates from underrepresented groups for reasons similar to the rationale for the more direct approach

3. For examples, see Green (2017), Emerson (2017), Toten (2017), and *CEB Talent Daily* (2017).

discussed above. Third, if it breaks stereotypes about skillsets needed in an industry or occupation, this may make some groups more willing to opt in (similar to studies showing increased selection into competitive settings by women when altering gender-task stereotypes—see, for example, Shurchkov 2012; Flory, Leibbrandt, and List 2014). Breaking stereotypes may also push individuals across a System 1/System 2 divide (Stanovich and West 2000; Kahneman 2011), replacing an intuitive aversion to non-diverse environments with more reflective responses that cause some to opt in as a payoff-maximizing choice. This leads to our second main hypothesis.

Hypothesis 2: Projecting a desire for employees from a variety of academic fields and training, including those not commonly associated with the industry, will increase interest by underrepresented groups and achieve greater demographic diversity among selected candidates.

Finally, it is also important to know whether words are enough, or if a proven track record or commitment to employee diversity is important. The practical relevance of this question is particularly strong for organizations with little existing diversity. There are good reasons to expect facts of this nature may matter. On the one hand, individuals might interpret diversity-friendly statements as cheap talk—a cynical effort by the firm to give the appearance of trying to address a public demand or help promote a social good but lacking any real commitment behind it. On the other hand, even if the appeal for diversity is interpreted as being made in good faith, it might have far less effect (and potentially even backfire) if not backed up with evidence regarding the firm’s priorities, values, or current conditions. For example, an attempt to get underrepresented groups to apply may be interpreted by some as an indicator of a “diversity problem” (little current diversity, lack of support once hired, etc.), which could dampen the impacts of a diversity appeal (or even push away some candidates that might have otherwise applied). This leads to our third main hypothesis.

Hypothesis 3: Including factual information to support claims about the value of diversity among employees to the organization will increase the ability of pro-diversity statements to raise interest among underrepresented groups.

B. Experiment Design Overview

To test these hypotheses and uncover how to increase workplace diversity, we run a field experiment in a high-profile sector perceived as lacking in diversity. The financial services industry is widely thought of as dominated by white males—this is discussed, for example, in the management and sociology literatures (Rivera 2015) and regularly surfaces in the popular press.⁴ This may be for good reason. In 2010, African Americans accounted for 2.7 percent of senior staff in financial services, while Hispanics accounted for 2.9 percent (Government Accountability Office 2013). In 2014, a survey by the *Financial Times* found 23.7 percent of senior roles in finance were occupied by

4. A sample of headlines from the last three years includes, for example, “‘White Male’ Culture at Banks Is Difficult Even for One of the Industry’s Top Leaders,” “Wall Street’s Young Bankers Are Still Mostly White and Male, Report Says,” “What It’s Like to Be Something Other Than White and Male in the Hedge Fund Business,” “The Financial Industry Doesn’t Want You to Know about Its Lack of Diversity,” “These Charts Show Just How White and Male Wall Street Really Is.” Also see Brymer (2016), Crowe and Kiersz (2015), Alden (2014), and *Financial Times* (2017).

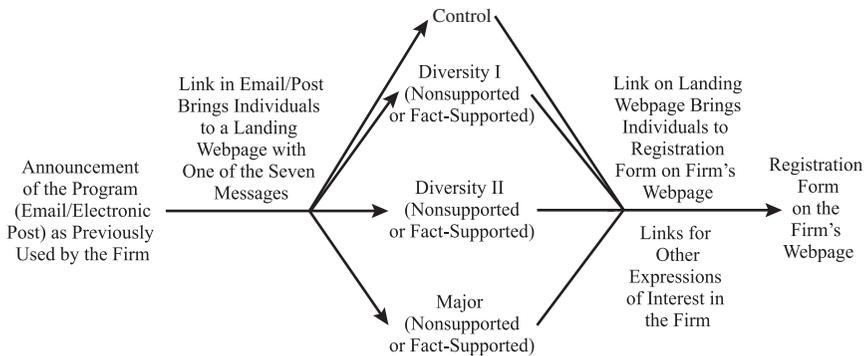


Figure 1
Overview of Experimental Protocol

Notes: Treatments are randomly assigned at individual level.

women, while a recent study of U.S. financial services firms put the percentage for women even lower—at 20–22 percent.⁵ Perceptions of lack of diversity therefore appear to have at least some support in the data. We take advantage of this setting to test whether different types of information can affect interest by underrepresented groups and recruiting outcomes.

The experiment is embedded in a recruitment drive to fill positions in a professional development program. This highly selective program introduces undergraduate freshmen and sophomores to careers in the financial industry and helps them build their professional networks and strengthen important skills like interviewing, elevator-pitching, résumé design, etc. Acceptance provides a critical “foot in the door” opportunity for participants, substantially boosting their chances of post-graduate employment at the firm and their readiness to enter and advance in the financial sector more broadly.⁶

Figure 1 illustrates the experimental protocol. The firm recruits participants by sending an email advertisement to its nationwide network of campus contacts, including career services centers, student organizations, individual business schools, departments, and job boards. The email (shown in [Online Appendix Figure A1](#)) consists of a brief announcement, along with a hyperlink to a webpage that has information about the firm, the program, and how to apply. Individuals who click on the hyperlink in the announcement first enter their name and university, are then randomized into one of the treatments, and then proceed to the landing page where the treatment is delivered. More precisely, we use the landing page to test the effects of different types of information on increasing interest from underrepresented groups. A randomly assigned statement at the top of the page either (i) provides information about the value the

5. See *Financial Times* (2017, <https://fig.ft.com/managements-missing-women-data/>) and Jaekel and St-Onge (2016).

6. In 2016, participants in this program were four times more likely than other applicants to be hired as an intern, the majority of whom become full-time employees after graduation. Since the program is also designed to build the participants’ professional image, acumen, and network in the financial sector more broadly, it is also likely to increase entrance to other major financial firms.

Table 1
Experimental Design

Message Type	Treatment	Support Type	Description
Control	T0	N/A	A diversity-neutral statement
Diversity I (direct diversity signal)	T1	Nonsupported	Statement emphasizing firm's valuation of diversity as a competitive advantage that raises firm productivity
	T2	Fact-supported	T1 plus listing of exact percentages of women and ethnically diverse new hires in the prior year
Diversity II (direct diversity signal)	T3	Nonsupported	Statement emphasizing firm's valuation of diversity as a key part of the company's culture
	T4	Fact-supported	T3, presented as quote spoken by the company's CEO, proving endorsement at the highest levels of company leadership
Major (indirect diversity signal)	T5	Nonsupported	Statement encouraging individuals from a variety of fields of study
	T6	Fact-supported	T5 plus listing the exact number of majors represented in last year's program

Notes: Individuals randomly received one of seven different treatment statements. The experiment follows a three by two design. There are three main message types (Diversity I, Diversity II, Major), and for each there are two support types (fact-supported and nonsupported). For the exact wording of each treatment, see [Online Appendix Table A1](#).

firm places on diversity among its employees, (ii) encourages individuals from a broad range of fields of study to apply, or (iii) includes a neutral statement to serve as our control condition. Statements were randomly assigned at the individual level.⁷ Besides the treatment script, the webpage is identical for all individuals who visit it. (See [Online Appendix Figure A2](#).)

C. Treatments

The treatments are designed to identify best practice and to test the three hypotheses outlined above, plus a minor hypothesis on whether it matters if the reason for wanting a diverse group of employees is its impact on firm performance or instead its inherent value to the firm's culture. Table 1 summarizes the treatment messages.⁸ (For the complete text

7. We removed a few students from the sample who entered the landing page multiple times and were exposed to different messages.

8. Besides the treatment language, the webpage for each treatment is identical.

in each treatment, see [Online Appendix Table A1.](#)) Our experiment follows a three-by-two design, with three main message types, each split into a fact-supported and a non-supported version. The first two message types signal the value of diversity to the firm: *Diversity I* emphasizes its value to firm productivity, and *Diversity II* emphasizes its value to firm culture. These signals are direct, communicating a desire by the firm to employ people from a variety of backgrounds and a high value placed on diversity per se. These treatments enable us to test Hypothesis 1.

The third message type tests an indirect approach to see if a more diverse pool can be attracted without explicitly appealing to applicant origins or social background, or even mentioning diversity per se. This message type (*Major*) attempts to encourage interest from ethnic minorities and women by appealing to a broader variety in major field of study than is typically associated with finance careers, citing subject areas with greater representation by ethnic minorities and women—such as ethnic studies, nursing, and psychology.⁹ This treatment is aimed at examining Hypothesis 2.

Finally, we split each of the three message types into one treatment with no evidence to support its broader claim (*nonsupported*) and one that presents evidence to back up the statement (*supported*). The facts used are the high percentage of recent hires accounted for by ethnic minorities and women (*Diversity I*), a direct quote from the company's CEO showing endorsement of the statement at the highest levels of company leadership (*Diversity II*), and information on breadth of student majors among past successful applicants (*Major*). This treatment allows us to examine Hypothesis 3.

At the bottom of the webpage, after seeing one of the seven messages, individuals can click on a hyperlink to learn more about the program and submit their application. By linking with the firm's applicant tracking database, we are able to capture the behavior of each individual who lands on the webpage—which treatment message she sees, whether she expresses interest in the program by starting an application or asking to be notified of similar events, whether she submits the application, and whether she is selected to participate in the program. Employees at the firm selecting participants from the pool of candidates do not know which treatment applicants received.

D. Outcome Variables

We examine three critical recruiting outcomes to analyze the impact of the different statements on race and gender diversity: (i) the percentage of different types of individuals who express interest by beginning an application or asking to be notified of similar future events, (ii) the percentage of different types of individuals who complete and submit an application to the program, and (iii) the percentage of different types of individuals selected for the program. We also look at two measures for each variable that are key to assessing impacts on diversity outcomes: the gap between underrepresented groups and overrepresented groups, and the behavior of each group (underrepresented and overrepresented) considered on its own.

9. In 2015, for example, the proportion of majors accounted for by African and Hispanic Americans rises 62 percent in moving from business to ethnic studies and rises 25 percent in moving from business to psychology (13 and 7 percentage points, respectively), while the proportion of majors accounted for by women rises 64 percent in moving from business to psychology and 51 percent in moving from business to ethnic studies (30 and 24 percentage points)(National Center for Education Statistics 2018).

The signal of interest provides the most complete measure of treatment effects of the statements on the appeal of the opportunity to individuals, since the decision to submit an application conditional on interest is likely driven by specifics of the program such as its date or location.¹⁰ However, in addition to applicant interest, submitted applications are perhaps just as important to organizations and to diversity recruiting efforts, so we report the effects on both. We also examine evidence on whether the impacts on interest among underrepresented/minority groups translate to impacts on their representation among selected candidates. This last measure is an indicator for the qualifications of the marginal candidates attracted by the treatments, since the selection committee was blind to the treatments.

E. Definition of Underrepresented Groups and Identification of Diversity Characteristics

In this study, we place particular emphasis on certain underrepresented races. In the work setting and related entry-level positions in the financial industry that we examine, the two groups that are by far the most underrepresented and highest priority for the firm's diversity recruiting objectives are African Americans and Hispanics.¹¹ While women and Asians are also underrepresented in financial industry leadership roles at the upper levels, it is not clear they are underrepresented in entry-level up to mid-level leadership positions in the financial industry.¹²

To identify ethnicity/race (and gender), we use two independent sets of data. First, we use data from self-reports. Individuals had the possibility to self-identify their ethnicity/race (and gender) during the application process, and 285 individuals did so. Second, we employed an independent research assistant to code the ethnicity. To do so, the research assistant found each individual on social networking sites using their name and university. Ethnicity was determined based on information such as native language, school clubs and societies of which they were a member, hometown, profile pictures, etc. When ethnicity was not immediately clear from this information, other publicly available personal websites, blogs, or news articles were used. At times, research into the etymology of names was also considered to help determine ethnicity. We use the same

10. For example, some percentage of those who would otherwise apply will learn they have conflicts with the date of the program or decide it is located too far away and thus ultimately will not complete and submit their application.

11. Other significantly underrepresented (but much smaller) groups are, for example, Native American and Hawaiian. However, due to their very small numbers, we do not separately analyze their behavior and instead pool them in the group "other" ethnicities, together with Asians, white, two or more races, and nonidentified. The vast majority of this group (92 percent) is composed of whites and Asians, which are not underrepresented in the finance sector. The results for whites and Asians only are very similar. See [Online Appendix Tables A2, A5, and A6](#) for further details on ethnic composition of candidates.

12. In 2015 Asian Americans accounted for about 5 percent of the U.S. labor force (Bureau of Labor Statistics: <https://www.bls.gov/opub/reports/race-and-ethnicity/2015/home.htm>), 11 percent of financial industry professionals, and 8 percent of first- and mid-level management (EEO: <https://www1.eeoc.gov/eeoc/statistics/employment/jobpat-eeo1/2015/index.cfm#centercol>; https://www1.eeoc.gov/eeoc/statistics/employment/jobpat-eeo1/2015/index.cfm#select_label); while women accounted for about 47 percent of the US labor force, (U.S. Department of Labor: https://www.dol.gov/wb/stats/NEWTATS/facts/women_If.htm#one), 52 percent of financial industry professionals, and 48 percent of first- and mid-level management (same EEO references as above). Women make up approximately 45 percent of the S&P 500 companies' workforce (Catalyst 2015), and Asians hold about 47 percent of professional jobs in Silicon Valley tech companies and roughly a quarter of finance sector professional employment (U.S. Equal Employment Opportunity Commission 2015).

categories for the variables ethnicity/race as used by the company in its application form. We coded ethnicity as follows: Asian; Black or African American; Hispanic or Latino; White; Two or More Races; Native, Hawaiian, or Other; and Cannot Tell.

As a check on the accuracy of the second approach in identifying race (and gender), we compare the race data identified through this process to the 285 individuals who applied and for whom we already had self-identified race. We find that all individuals who we identified as African American or Hispanic through this approach also self-identified as such or as having two or more races, that more than 99 percent of those identified through this process as female also self-identified as female, and that more than 98 percent of those identified as male also self-identified as male—this suggests that we were able to quite accurately identify ethnicity and gender.¹³

III. Experimental Findings

A. Overview and Global Effects

In total, the experiment generated a relatively large sample of 1,121 individuals, with a substantial number belonging to the underrepresented group ($N=166$, 14.8 percent of the sample: 6.3 percent African Americans, 8.5 percent Hispanic).¹⁴ As expected, Asians (47.7 percent) and women (48.4 percent) are not underrepresented in our sample.¹⁵ Of the 1,121 individuals, 385 signal interest after viewing the treatment message and program information (62 African Americans and Hispanics and 211 women), 285 submit applications (50 African Americans and Hispanics and 156 women), and 47 are selected to participate in the program (15 African Americans and Hispanics and 33 women). Randomization of individuals into the treatment cells resulted in a fairly well-balanced distribution—each message was seen by approximately 160 individuals, ranging from 150 individuals who saw the *Diversity I (supported)* message to 165 individuals who saw the *Diversity II (nonsupported)* message. About 15 percent of each treatment group is African American or Hispanic (ranging from 13 percent to 18 percent across the seven groups), and about 48 percent of each treatment group is women (from

13. We also find that the Spearman's correlation coefficient across the values for race is 0.80 ($p=0.000$) when looking across all ethnicities – including white, Asian, Native/Hawaiian/Other – and is 0.96 ($p=0.000$) for gender (see [Online Appendix Tables A5 and A6](#)). Note that any measurement error in the identification of race and gender (that is, noise) makes it *more* difficult to identify the impact of treatments targeted to affect racial/gender minorities. Throughout the analysis, we use the coded information. Given the high correlation between coded and self-identified ethnicity and gender, we are confident that we capture the treatment effect on ethnic minorities and women.

14. While our sample size of subjects from underrepresented groups may appear somewhat limited at first glance, we note that 166 individuals is relatively large compared to total sample sizes in many typical laboratory experimental studies. One of the advantages of lab experiments is that they draw from a population with less variation than is common in field experiments (university undergraduates) and are thus able to obtain relatively precise estimates even with small samples. In this field experiment, we also fortunately benefit from this advantage, as our sample is drawn from undergraduates.

15. Overall, 1,264 individuals clicked on the hyperlink in the email announcement, entered their name and university, and proceeded to the landing page where the treatment was delivered (33 individuals clicked on the hyperlink but did not pass through all the way to the landing page). However, we had to exclude 135 individuals from the sample who logged in multiple times from different IP addresses and saw different treatment messages and eight individuals who saw the message after they had already applied to the program. This leaves us with a total sample of 1,121 individuals who saw the control or one of the six treatment messages.

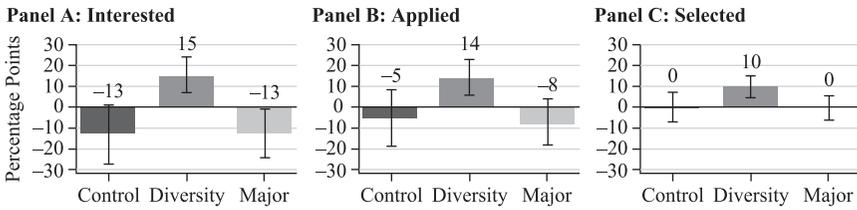


Figure 2
Impact of Messages on Ethnicity Gap

Notes: Difference between percentage of underrepresented and nonunderrepresented individuals expressing interest (Panel A), submitting an application (Panel B), and being selected (Panel C). 90 percent confidence intervals are calculated with robust standard errors.

41 percent to 53 percent), with no significant differences in percentages across treatment groups. (See also [Online Appendix Table A2.](#))

We first examine impacts of the statements on interest of all individuals, ignoring demographic characteristics. We find that interest in the company and/or a career in the financial sector is higher in all six treatments compared to the control. Pooling all six treatment messages together, we see they raise the overall proportion of individuals interested in the program by 25 percent (seven percentage points), from 28 percent of those who visited the landing page with the status quo *Control* message, to 35 percent of those who visited the landing page with one of the six messages communicating employer prioritization of diversity among its employees (χ^2 -test, $p=0.057$). Splitting by the two different types of approaches to attracting employee diversity (*Diversity* vs. *Major*), we find they each lead to similar positive impacts on individuals overall—a seven percentage point rise for *Diversity* (χ^2 -test, $p=0.076$) and an eight percentage point rise for the *Major* messages (χ^2 -test, $p=0.069$).

B. Effects on Underrepresented Groups

We now turn to impacts of the statements on the underrepresented groups in recruiting outcomes. There are two distinct and important dimensions with respect to treatment effects on demographic composition of candidates and recruits: impacts on the *gap* between underrepresented and nonunderrepresented and impacts on *individuals* from underrepresented and nonunderrepresented groups. We examine both. We also first pool the *Diversity* messages separately from the *Major* messages since they represent two fundamentally different approaches to raising interest among underrepresented groups and because we find little difference in impacts within the *Diversity* category overall (see further below).

Figure 2 illustrates the gap between underrepresented and nonunderrepresented ethnicities in the proportion of individuals who express interest (Panel A), the proportion who submit an application (Panel B), and the proportion who are selected (Panel C) in each of the three main message types (*Control*, *Diversity*, and *Major*). We see that in the *Control* condition, underrepresented ethnicities were 13 percentage points less likely to express interest than nonunderrepresented, five percentage points less likely to apply, and did not differ in the rate that they were selected. Turning to

the *Diversity* condition, the picture sharply changes: individuals from underrepresented groups are now 15 percentage points *more* likely to be interested, 14 percentage points *more* likely to apply, and ten percentage points more likely to be selected. That is, the *Diversity* condition changes the gap by 28 percentage points ($p=0.006$; OLS with robust SE) in favor of underrepresented groups for expressing interest, by 19 percentage points ($p=0.052$) for applying (in both cases, not only closing, but reversing the gap), and by ten percentage points ($p=0.073$) for being selected.

Table 2 reports results from three linear probability models that examine the robustness of these findings, accounting for gender of the candidate. Panel A shows coefficient estimates for regressions of the outcomes on gender, ethnicity, treatment type, as well as ethnicity-by-treatment and gender-by-treatment interactions. The coefficient estimates for the interaction term *African/Hispanic American* \times *Diversity* represent the estimated impacts of the *Diversity* treatment on the ethnicity gap for each outcome variable (that is, the difference-in-difference estimates – the amount that *Diversity* increases each outcome variable for underrepresented ethnicities, relative to nonunderrepresented). As the estimates show, *Diversity* raises the percentage of underrepresented groups expressing interest by 30 percentage points ($p < 0.01$), the percentage applying by 21 percentage points ($p < 0.05$), and the percentage being selected by 11 percentage points ($p = 0.053$), above that of nonunderrepresented.

Turning to how the treatments impact behavior and outcomes for individuals, Figure 3 illustrates the effects among the two different groups (African American/Hispanic and other ethnicities). It also sheds light on what is driving the above effects on the ethnicity gap for each outcome. The top row reports the percentage of African Americans and Hispanics who express interest (Panel A), submit applications (Panel B), and are selected (Panel C) depending on the three main message types in the job advertisement (*Control*, *Diversity*, *Major*).

As we see in the Panels A and B, African Americans and Hispanics strongly respond to the *Diversity* messages. Panel A shows that their interest in the open position almost triples in moving from the *Control* to the *Diversity* message. While only 17 percent of ethnic minorities signal interest in *Control*, 48 percent do so in *Diversity* (χ^2 , $p = 0.005$). Further, in Panel B we observe that the proportion of ethnic minorities submitting an application more than doubles in moving from *Control* to *Diversity*, going from 17 percent to 38 percent (χ^2 , $p = 0.05$).

The encouragement effect of *Diversity* on African and Hispanic Americans also translates into a threefold increase in the likelihood of being selected, compared to their likelihood of being selected in *Control*, and induces a sixfold increase in the likelihood of being selected relative to nonunderrepresented individuals. This provides indirect evidence that the *Diversity* treatments not only encourage applications from African Americans and Hispanic Americans per se, but from highly qualified African Americans and Hispanic Americans, as the selection committee is blind to treatment and does not know which message type each applicant sees. More precisely, while only 4 percent of underrepresented ethnicities in *Control* are selected for openings, 12 percent of ethnic minorities in *Diversity* are selected (χ^2 , $p = 0.25$).¹⁶ Furthermore, while the same

16. The share of selected applicants is another way of comparing qualifications of applicants across conditions. If the *Diversity* messages attract highly qualified African Americans/Hispanics, the share of selected African American/Hispanic applicants in *Diversity* should be at least as large as in *Control* (and *Major*). This is indeed the case: while 33 percent of African American/Hispanic applicants are selected in *Diversity*, the shares

Table 2
Probability of Expressing Interest, Submitting Application, and Being Chosen

	Interested (1)	Applied (2)	Selected (3)
Panel A: Treatment Effects on Ethnicity Gap, Gender Gap, and Overall			
Diversity	0.0336 (0.0539)	0.0222 (0.0480)	-0.0109 (0.0198)
Major	0.0892 (0.0613)	0.0674 (0.0550)	-0.0130 (0.0199)
African American/Hispanic	-0.145 (0.0884)	-0.0614 (0.0863)	-0.00758 (0.0436)
African American/Hispanic × Diversity	0.301*** (0.104)	0.207** (0.101)	0.107* (0.0552)
African American/Hispanic × Major	0.0271 (0.114)	-0.00641 (0.109)	0.0155 (0.0550)
Female	0.106 (0.0713)	0.0798 (0.0655)	0.0430 (0.0327)
Female × Diversity	-0.0128 (0.0806)	-0.0141 (0.0740)	-0.0207 (0.0361)
Female × Major	-0.0312 (0.0890)	-0.0191 (0.0819)	0.0252 (0.0398)
Constant	0.249*** (0.0472)	0.182*** (0.0421)	0.0241 (0.0179)
Panel B: Treatment Effects on Ethnic-Minority Individuals			
Diversity Effect on African American/Hispanic Men [<i>F</i> -test <i>p</i> -value]	0.335*** [0.0014]	0.229** [0.0252]	0.096* [0.0607]
Diversity Effect on African American/Hispanic Women [<i>F</i> -test <i>p</i> -value]	0.322*** [0.0012]	0.215** [0.0262]	0.0756 [0.1941]
Major Effect on African American/Hispanic Men [<i>F</i> -test <i>p</i> -value]	0.1163 [0.3092]	0.0610 [0.5799]	0.00252 [0.9601]
Major Effect on African American/Hispanic Women [<i>F</i> -test <i>p</i> -value]	0.0851 [0.4381]	0.0419 [0.6898]	0.0277 [0.6359]
Observations	1,121	1,121	1,121

Notes: Estimates from a linear probability model. Dependent variables are dummy variables taking the value one if the individual expresses interest in the program (Model 1), submits an application (Model 2), or is selected (Model 3). The explanatory variables Diversity, Major, African American/Hispanic, and Female are dummy variables taking the value one if the individual is in the respective treatment or demographic group. Diversity Effect on African American/Hispanic Men is the sum of the coefficients for Diversity and African American/Hispanic × Diversity, and Diversity Effect on African American/Hispanic Women is the sum of the same coefficients plus the coefficient for Female × Diversity. Major Effect on African American/Hispanic Men is the sum of the coefficients for Major and African American/Hispanic × Major, and Major Effect on African American/Hispanic Women is the sum of the same coefficients plus the coefficient for Female × Major. Robust standard errors in parentheses, *F*-test *p*-values in brackets; **p*<0.1, ***p*<0.05, ****p*<0.01.

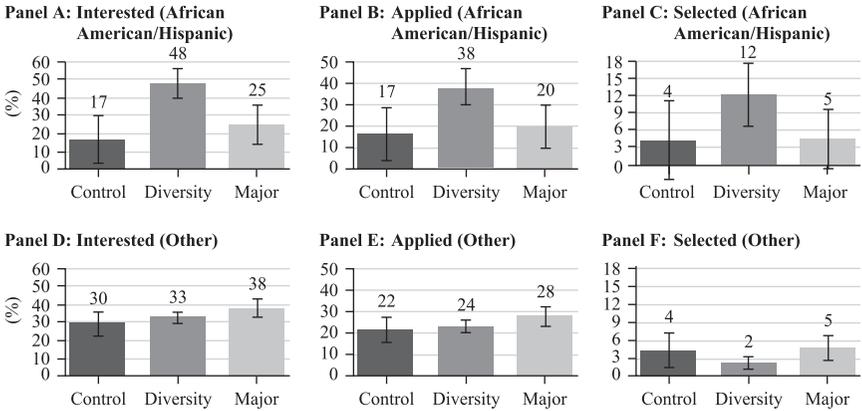


Figure 3
Impact of Messages

Notes: Percentage of individuals expressing interest (Panels A and D), submitting an application (Panels B and E), and being selected (Panels C and F). The top (bottom) row shows the respective percentages for African Americans and Hispanics (other individuals). 90 percent confidence intervals are calculated with robust standard errors.

proportion of underrepresented ethnicities as nonunderrepresented are selected for the program under *Control* (4 percent), under the *Diversity* message type underrepresented ethnicities are ten percentage points (500 percent) more likely to be selected than nonunderrepresented ($\chi^2, p=0.000$).

The estimates in Panel B of Table 2 confirm the positive and highly significant impact of *Diversity* messages on underrepresented ethnicities’ interest (Column 1, $p=0.001$ for men, $p=0.001$ for women) and application rates (Column 2, $p=0.025$ for men, $p=0.026$ for women), compared to interest and application rates of underrepresented ethnicities in the *Control* condition. Furthermore, Column 3 shows that for individuals from underrepresented groups who view the page, the *Diversity* messages raise the likelihood that they will be selected for the program by an estimated 9.6 percentage points ($p=0.061$) for males and an estimated 7.5 percentage points ($p=0.194$) for females, compared to their likelihood of being selected under the *Control*.

To summarize, nonparametric tests and regression estimates both show substantial effects from the *Diversity* messages on all three recruiting outcomes—interest, applications, and selection. This is true both when looking at the impacts on the difference between underrepresented and nonunderrepresented ethnicities, as well as the effects on underrepresented ethnicities alone (ignoring the effects on nonunderrepresented). Indeed, impacts on the gap are almost entirely driven by the effects on underrepresented group individuals. Finally, we find no significant differences in treatment effects between African Americans and Hispanics ($p>0.43$).¹⁷

are 17 percent and 40 percent in *Control* and *Major*, respectively (and 36 percent among African American/Hispanic applicants who do not see any of the treatment messages).

17. In attracting candidate interest, for example, the *Diversity* statements raise interest among African Americans by 40 percentage points for men ($p<0.01$) and 37 percentage points for women ($p<0.01$), and raise interest among Hispanic Americans by 29 percentage points for men ($p<0.05$) and 28 percentage points for

Several other notable findings are illustrated in Figures 2 and 3. First, we observe in Figure 2 that the *Major* treatments appear to have no positive impact on the gap between underrepresented ethnicities and nonunderrepresented for any of the three outcomes, a finding confirmed in the regression results in Panel A of Table 2, where the estimates for the interaction term *African/Hispanic American* \times *Major* are quite small and not significantly different from zero. Turning to Figure 3 helps us see why—the top row shows underrepresented ethnicities do not respond nearly as strongly to the indirect *Major* treatments. While their expression of interest and submission of applications are both higher in *Major* than in *Control* (increases of eight percentage points and three percentage points, respectively), the difference is not statistically significant (χ^2 -tests, $p=0.428$, and $p=0.704$). This is confirmed by the results shown in Panel B of Table 2, where the estimated impacts of *Major* are positive for all three outcomes, but never significantly so (regression estimate *F*-tests, $p>0.309$).

Figures 2 and 3 further illustrate the strong effect of the direct *Diversity* messages compared to the indirect *Major* messages. As pointed out earlier, the ethnicity gap in rates of interest and applications reverses in *Diversity*, which is not the case in *Major*. In addition, while 48 percent of African Americans and Hispanics signal interest upon seeing a *Diversity* message, only 25 percent do so in *Major* (χ^2 -test, $p=0.010$), and while 38 percent of African Americans and Hispanics apply under *Diversity*, only 20 percent do so under *Major* (χ^2 -test, $p=0.042$). *F*-tests for the regressions in Table 2 (not shown) confirm the same pattern ($p=0.014$ and $p=0.046$ for African American/Hispanic men, $p=0.010$ and $p=0.042$ for African American/Hispanic women). The difference between the impact of *Diversity* and the impact of *Major* on the selection rates is in the same direction (stronger impacts for *Diversity*), but is less pronounced (χ^2 -tests, $p=0.150$; regression estimate *F*-tests, $p=0.043$ and $p=0.338$ for African American/Hispanic men and women, respectively).

Another noteworthy finding relates to the response to diversity appeals among individuals from ethnic groups that are not underrepresented in this sector. In the bottom row of Figure 3, we see that there is no discouragement effect on rates of interest or applications from the *Diversity* messages on overrepresented ethnic groups (whites and Asians in this setting). Ethnic nonminority candidates, upon seeing the *Diversity* and *Major* messages, are even slightly more likely to signal interest (33 percent in *Diversity* and 38 percent in *Major*, compared to 30 percent in *Control*) and to apply (24 percent in *Diversity* and 28 percent in *Major*, compared to 22 percent in *Control*), although these differences are not significant (χ^2 -tests, $p=0.483$ and $p=0.651$ for *Diversity*, $p=0.105$ and $p=0.166$ for *Major*). This is also visible in the regression estimates in Panel B of Table 2, which show positive and nonsignificant coefficients for the impacts of *Diversity* and *Major* on interest and application rates among both men and women in the nonunderrepresented group. This provides evidence that the *Diversity* treatments benefit members of underrepresented groups and helps employer diversity objectives, without discouraging members of overrepresented groups. Combined with the findings discussed earlier on the positive overall impacts on interest across all individuals, this suggests little to no downside in recruiting outcomes from using the diversity appeals.

Figure 4 shows the patterns when considering the *Diversity I* and *Diversity II* treatments separately. Overall, the difference in impacts is negligible for both underrepresented

women ($p<0.05$), the difference in effects between African Americans and Hispanics not significant ($p>0.50$). See [Online Appendix 1](#) and [Table A7](#) for full results on each ethnic minority group examined separately.

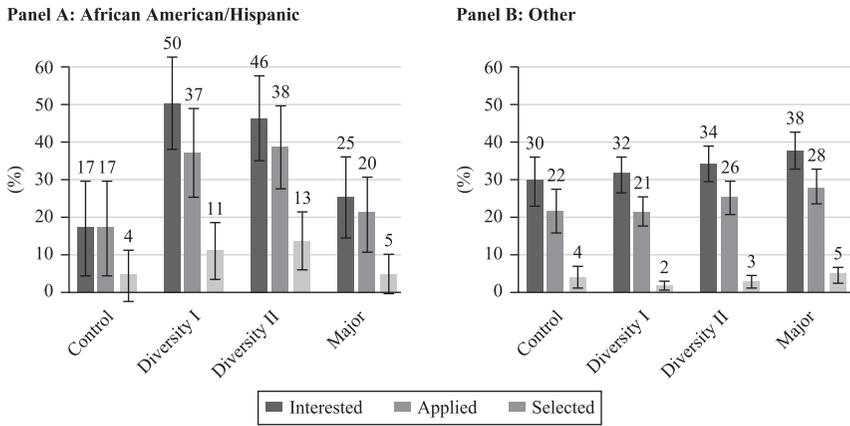


Figure 4

Proportion of Individuals Expressing Interest, Applying, and Being Selected

Notes: Percentages are shown for African Americans and Hispanics (Panel A) and for other individuals (Panel B). 90 percent confidence intervals are calculated with robust standard errors.

ethnicities (Panel A) and for individuals outside the underrepresented groups. Similarly, impacts of each treatment on the ethnicity gap are also quite close to each other—moving the gap in favor of underrepresented ethnicity by 31 versus 25 percentage points for interest, by 21 versus 17 percentage points for completed applications, and by nine versus ten percentage points for being selected. Regression results reported in [Online Appendix Table A3](#) confirm that there are no significant differences in the effects on candidates from underrepresented ethnicities from *Diversity I* compared to *Diversity II*, and that the same results discussed above for *Diversity* hold when considering the impacts of the two statements separately. We therefore use the pooled *Diversity* treatments as the basis for our main findings, as discussed above.

C. Gender Findings

Although women are not underrepresented in our setting, they may still respond to our treatments. Turning our analysis to women, we find they appear much less sensitive in general to the treatments as compared to underrepresented ethnicities. While their interest and application rates are higher in *Diversity* and *Major* than in *Control*, these differences are not significant at conventional levels. For example, the proportions of women who signal interest are 33 percent in *Control*, 40 percent in *Diversity*, and 40 percent in *Major*, while the proportions who complete an application are 25 percent in *Control*, 29 percent in *Diversity*, and 30 percent in *Major* ($p > 0.267$ for pairwise χ^2 -tests for all four treatment–control comparisons). For applicant selection, 7 percent of all women who see the *Control* statement are selected compared to 5 percent of those who see a *Diversity* statement and 8 percent of those who see a *Major* statement. There is also little evidence of effects of the treatment statements on the gender gap. As shown in Table 2, the estimates for the gender-by-treatment interaction terms are small and not significantly different from zero.

This picture changes little when considering the impacts of *Diversity I* and *Diversity II* separately. When looking at the comparisons between *Control* and *Diversity I*, and *Control* and *Diversity II*, interest and application rates among women, as well as selection by the firm, are once again generally higher than in the *Control*, but none of the differences are statistically significant at conventional levels (χ^2 -tests, $p > 0.21$). In addition, neither the *Diversity I* nor the *Diversity II* treatment has a statistically significant impact on the gender gap across the three outcomes, as shown by the coefficient estimates for the gender-by-treatment interaction terms in [Online Appendix Table A3](#).

D. Factual Support

The third main hypothesis this experiment was designed to test is that adding concrete facts or statistics to diversity or inclusivity statements increases their impact by raising their credibility. Interestingly, we find little evidence that providing factual support for the pro-diversity orientation of the firm has an effect. Pro-diversity statements without supporting evidence were no less effective than those with supporting evidence, whether facts about diversity among recent recruiting outcomes were presented as in *Diversity I* or *Major* or revealed as a quote by the company's CEO as in *Diversity II*. Pairwise tests of the difference between the fact-supported and nonsupported statements within *Diversity I*, *Diversity II*, and *Major* show no significant difference for signaling interest (underrepresented ethnicities: $p > 0.73$; women: $p > 0.21$; all individuals: $p > 0.29$), for applying (underrepresented ethnicities: $p > 0.29$; women: $p > 0.26$; all individuals: $p > 0.24$), and for being selected (underrepresented ethnicities: $p > 0.17$; all individuals: $p > 0.39$). The only exception is among women for the selection outcome, where a higher proportion of women who saw the fact-supported version of *Diversity I* were selected for the program compared to those who saw the nonsupported version (6.35 percent compared to 0 percent, χ^2 -test, $p = 0.02$). (The difference in selection rates for women between fact-supported and nonsupported under *Diversity II* and *Major* are not significant, $p > 0.25$.) Linear probability model regression results reported in [Online Appendix Table A4](#) confirm these results.

One potential explanation for the lack of observed impact of fact-supported statements is that their impact may depend on the position advertised. For example, the CEO statement, provided in Treatment 4, may be less important for students for a one-day professional development program than for job-seekers for full-time, high-profile positions. Another potential explanation is that the prospective applicants did not perceive the additional information to be more informative than the statement itself. For example, listing the share of diverse new hires in Treatment 2 may not help if the share is about equal to what the candidates expect—this new information may not change their prior beliefs. While our current experimental design does not allow us to disentangle competing explanations for the ineffectiveness of the fact-supported statements, this provides an important direction for future research.

IV. Discussion and Conclusion

Companies have tried different strategies to increase diversity, and there is some evidence that affirmative action, such as quotas (Balafoutas and Sutter 2012; Niederle and Vesterlund 2013; Ibanez and Riener 2018), mentoring and advice

institutions (Blau et al. 2010; Rodriguez-Planas 2012; Brandts, Groenert, and Rott 2015) or changes in remuneration structure (Niederle and Vesterlund 2007; Flory, Leibbrandt, and List 2014), can influence job-seeker sorting in the labor market, with impacts on gender and ethnic diversity in the workplace. The findings we report highlight the value of a far less invasive approach, using simple adjustments in language to signal important dimensions of a firm's value of employee diversity. Our results suggest that signals valuing workplace diversity have important implications for net impacts on the size and demographic composition of applicant pools, and more generally for job-seekers heading into high-profile careers.

The implications of our results are important in several regards. On the one hand, this method of using statements to signal active valorization of diversity offers employers and organizations an easy and cost-effective tool to enhance diversity in situations where using quotas or altering compensation structures is not feasible, costly to implement, or undesirable for other reasons. The fact that diversity messages increase the proportion of ethnic-minority applicants who are actually selected suggests the messages not only pull in more ethnic minority applicants, but pull in those who are strong candidates. On the other hand, identifying a method not requiring resource-intensive interventions and not linked with financial incentives has tremendous value. It points toward a range of potential applications far beyond firms and labor markets: since our results stem from simple changes in language and signals of values, they may speak to a broad range of environments where lack of diversity is a concern. Enrollment in higher education, political participation (voting, running for election), and civic engagement are just a few examples where language and signals might be leveraged to increase interest among underrepresented groups and break up monocultures (or prevent them from forming) in other spheres critical for the functioning of a healthy democratic society.

Finally, our finding that diversity messages sharply raise interest and application rates by African Americans and Hispanic Americans, but do not dissuade nonunderrepresented ethnicities is important. It suggests that this approach can lead to an increase in diversity without a dip in applicant quality or adverse effects on the majority group. In closing, it is important to bear in mind that our results highlight effective approaches to attract individuals from underrepresented minorities, but have less to say about how to retain them. When it comes to sustainably changing the demographic composition of employees in an organization, hiring is a critical piece, but forms part of a larger overall endeavor with other dimensions that must also be addressed in order to create lasting change.

References

- Alden, W. 2014. "Wall Street's Young Bankers Are Still Mostly White and Male, Report Says." *New York Times*, September 30.
- Apfelbaum, E., N. Stephens, and R. Reagans. 2016. "Beyond One-Size-Fits-All: Tailoring Diversity Approaches to the Representation of Social Groups." *Journal of Personality and Social Psychology* 111(4):547-66. <http://dx.doi.org/10.1037/pspi0000071>
- Ashraf, Nava, Oriana Bandiera, and Scott S. Lee. 2014. "Do-Gooders and Go-Getters: Career Incentives, Selection, and Performance in Public Service Delivery." STICERD Economic Organization and Public Policy Discussion Papers Series 54.

- Balafoutas, Loukas, and M. Sutter. 2012. "Affirmative Action Policies Promote Women and Do Not Harm Efficiency in the Laboratory." *Science* 335(6068):579–82.
- Bertrand, M., and E. Duflo. 2017. "Field Experiments on Discrimination." In *Handbook of Economic Field Experiments*, Volume 1, ed. Abhijit Vinayak Banerjee and Esther Duflo, 309–93. New York: Elsevier.
- Blau, F.D., J.M. Currie, R.T. Croson, and D.K. Ginther. 2010. "Can Mentoring Help Female Assistant Professors? Interim Results from a Randomized Trial." *American Economic Review* 100(2):348–52.
- Blau, F., and L. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." NBER Working Paper 21913. Cambridge, MA: NBER.
- Brandts, J., V. Groenert, and C. Rott. 2015. "The Impact of Advice on Women's and Men's Selection into Competition." *Management Science* 61(5):1018–35.
- Brymer, R.A. 2016. "What New Hiring Methods Say about Wall Street's Diversity Problem." *Fortune*, June 20.
- Buser, T., M. Niederle, and H. Oosterbeek. 2014. "Gender, Competitiveness, and Career Choices." *Quarterly Journal of Economics* 129(3):1409–47.
- Catalyst. 2005. "2005 Catalyst Member Benchmarking Report." <https://www.catalyst.org/wp-content/uploads/2019/02/2005MemberBenchmarkingReport.pdf> (accessed February 11, 2020).
- . 2015. "Pyramid: Women in S&P 500 Companies." <http://tinyurl.com/qcnwko5> (accessed February 11, 2020).
- CEB Talent Daily. 2017. "Diversity Backlash: Lean in Report Highlights America's Blind Spots on Gender Inequality." October 11.
- Crowe, Portia, and Andy Kiersz. 2015. "These Charts Show Just How White and Male Wall Street Really Is." *Business Insider*, August 25.
- Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi. 2013. "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service." *Quarterly Journal of Economics* 128(3):1169–218.
- Del Carpio, L., and M. Guadalupe. 2018. "More Women in Tech? Evidence from a Field Experiment Addressing Social Identity." Unpublished.
- Dobbin, F., and A. Kalev. 2016. "Why Diversity Programs Fail." *Harvard Business Review*, August 21.
- Emerson, Joelle. 2017. "Colorblind Diversity Efforts Don't Work." *Harvard Business Review*, September 11.
- Financial Times*. 2017. "Management's Missing Women Series." <https://www.ft.com/management-women> (accessed April 16, 2020).
- Flory, J.A., U. Gneezy, K.L. Leonard, and J.A. List. 2018. "Gender, Age, and Competition: A Disappearing Gap?" *Journal of Economic Behavior and Organization* 150:256–76.
- Flory, J.A., A. Leibbrandt, and J.A. List. 2014. "Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions." *Review of Economic Studies* 82(1):122–55.
- Gilbert, J.A., and J.M. Ivancevich. 2000. "Valuing Diversity: A Tale of Two Organizations." *Academy of Management Executive* 14(1):93–105.
- Giuliano, L., D. Levine, and J. Leonard. 2011. "Racial Bias in the Manager–Employee Relationship: An Analysis of Quits, Dismissals, and Promotions at a Large Retail Firm." *Journal of Human Resources* 46(1):26–52.
- Government Accountability Office. 2013. "Diversity Management: Trends and Practices in the Financial Services Industry and Agencies after the Recent Financial Crisis." <https://www.gao.gov/assets/660/653814.pdf> (accessed April 21, 2020).
- Green, Jeff. 2017. "Deloitte Thinks Diversity Groups Are Passé." *Bloomberg Businessweek*, July 19.

- Hansen, F. 2003. "Diversity's Business Case: Doesn't Add Up." *Workforce* 82(4):28–29.
- Hinton, I., J. Howell, E. Merwin, S. Stern, S. Turner, I. Williams, and M. Wilson. 2010. "The Educational Pipeline for Health Care Professionals: Understanding the Source of Racial Differences." *Journal of Human Resources* 45(1):116–56.
- Ibanez, M., and G. Riener. 2018. "Sorting through Affirmative Action: Three Field Experiments in Colombia." *Journal of Labor Economics* 36(2):437–78.
- Jaekel, Astrid, and Elizabeth St-Onge. 2016. "Why Women Aren't Making It to the Top of Financial Services Firms." *Harvard Business Review* <https://hbr.org/2016/10/why-women-arent-making-it-to-the-top-of-financial-services-firms> (accessed April 21, 2020).
- Kahneman D. 2011. *Decisions, Fast and Slow*. New York: Farrar, Straus, and Giroux.
- Kalev, A., F. Dobbin, and E. Kelly. 2006. Best Practices or Best Guesses? Assessing the Efficacy of Corporate Affirmative Action and Diversity Policies. *American Sociological Review* 71(4):589–617.
- Leibbrandt, A., and J. List. 2018. "When Equal Employment Opportunity Statements Backfire: Evidence from a Natural Field Experiment on Job-Entry Decisions." NBER Working Paper 25035. Cambridge, MA: NBER.
- Marinescu, Ioana, and Ronald Wolthoff. 2016. "Opening the Black Box of the Matching Function: The Power of Words." NBER Working Paper 22508. Cambridge, MA: NBER.
- McKay, P.F., and D.R. Avery. 2005. "Warning! Diversity Recruitment Could Backfire." *Journal of Management Inquiry* 14(4):330–36.
- National Center for Education Statistics. 2018. "Baccalaureate and Beyond." <https://nces.ed.gov/> (accessed February 11, 2020).
- Niederle, M., and L. Vesterlund. 2007. "Do Women Shy Away From Competition? Do Men Compete Too Much?" *Quarterly Journal of Economics* 122(3):1067–1101.
- Niederle, M., C. Segal, and L. Vesterlund. 2013. "How Costly Is Diversity? Affirmative Action in Light of Gender Differences in Competitiveness." *Management Science* 59(1):1–16.
- Preece, J., and O. Stoddard. 2015. "Why Women Don't Run: Experimental Evidence on Gender Differences in Competition Aversion." *Journal of Economic Behavior and Organization* 117:296–308.
- . 2016. "Run, Jane, Run! Gendered Responses to Political Party Recruitment." *Political Behavior* 38(3):561–77.
- Rivera, L. 2015. "Firms Are Wasting Millions Recruiting on Only a Few Campuses." *Harvard Business Review*. October 23, p. 1–8.
- Rodriguez-Planas, N. 2012. "Longer-Term Impacts of Mentoring, Educational Services, and Learning Incentives: Evidence from a Randomized Trial in the United States." *American Economic Journal: Applied Economics* 4(4):121–39.
- Shurchkov, Olga. 2012. "Under Pressure: Gender Differences in Output Quality and Quantity under Competition and Time Constraints." *Journal of the European Economic Association* 10(5):1189–213.
- Stanovich, K.E., and R.F. West. 2000. "Individual Differences in Reasoning: Implications for the Rationality Debate?" *Behavioral Brain Sciences* 23:645–65.
- Steele, C.M., and Joshua Aronson. 1995. "Stereotype Threat and the Intellectual Test Performance of African Americans." *Journal of Personality and Social Psychology* 69(5):797–811.
- Toten, M. 2017. "Backlash against Diversity and Inclusion Gains Momentum." *WorkplaceInfo*, May 3.
- U.S. Equal Employment Opportunity Commission. 2015. "Diversity in the Finance Industry Report." <https://www.eeoc.gov/eeoc/statistics/reports/finance/index.html> (accessed February 11, 2020).

Copyright of Journal of Human Resources is the property of University of Wisconsin Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.