PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES

Debiasing job ads by replacing masculine language increases gender diversity of applicant pools

Joyce C. He[a,1] and Sonia K. Kang[b]

Edited by Hazel Markus, Stanford University, Stanford, CA; received May 16, 2024; accepted December 23, 2024

Job advertisements for jobs in male-dominated fields tend to contain more masculine language, and a commonly proposed intervention to increase gender diversity in applicant pools is to remove this language. In our research, we offer predictions about the broader impact of such interventions on individuals who may not "fit" with traditional masculine identity. Across four multimethod studies (N = 37,920) spanning both field and lab settings, we demonstrate that a gender debiasing intervention that replaces masculine language in job postings with synonymous gender-neutral language increases application rates among women and men whose self-identities are less aligned with masculinity. Our research highlights that conceptualizing identity incongruence along a continuum can help explain when and why gender diversity initiatives foster the inclusion of a broader group of individuals than initially anticipated.

gender diversity | gender identification | belonging | intervention | masculinity

Significance

The language used in job advertisements often signals the "ideal" candidate profile and, therefore, who should apply. Unfortunately, job advertisements for male-dominated fields often use language that describes this ideal candidate as stereotypically masculine, implying that those with masculine identities are better suited for the role. We propose that reframing this language—by replacing stereotypically masculine words (i.e., words commonly associated with men) with synonymous gender-neutral words—can increase application rates. These increases are not only observed among women, but also among men who do not fit traditional masculine norms. We find evidence for these predictions across the lab and field, suggesting that efforts to make job advertisements more inclusive can benefit a broader range of individuals than initially anticipated.

Despite decades of intervention, women in Canada and the United States remain underrepresented in various areas of employment including STEM (1), certain business fields (e.g., finance; 2), and senior-level executive positions (3). Recent scholarship proposes "masculine defaults" as a cultural explanation for the maintenance of gender disparities in majority-male fields and occupations (4). These defaults reflect a cultural bias that prizes traits and behaviors associated with men and masculinity over those associated with women and femininity as normative, ideal, and valuable (5, 6). We focus on how masculine defaults can be signaled (and dismantled) in job advertisements, a critical recruitment and selection tool designed to attract preferred candidates (7, 8).

Job advertisements in majority-male contexts often communicate masculine defaults via stereotypically masculine language (e.g., competitive, ambitious, independent) (9). Accordingly, removing masculine language is frequently proposed as an intervention to increase gender diversity in these domains (e.g., 5). Previous interventions have often proposed and used a "feminizing" strategy (10), postulating that "...replacing the masculine wording [in job advertisements] with parallel feminine wording would increase women's interest in those jobs" (9, p. 120). However, "feminizing" has several limitations, including incitement of backlash and discrimination against counterstereotypical women who do not self-present as communal or feminine (11–13). Instead, we replace masculine language with gender-neutral language, an approach we call "gender debiasing." Rather than replacing masculine defaults with feminine defaults, this approach aims to "undo" gender altogether (14), e.g., replacing "firemen" with "firefighters" instead of "firemen and firewomen" (15). While we use the terms "gender debiasing" and "gender-neutral" throughout this manuscript, we recognize that achieving true "neutrality" by completely removing gender may be challenging or even impossible given the primacy of gender in shaping social cognition and perception (16, 17). Rather, our intervention aims to approximate neutrality by replacing overtly masculine language with relatively gender-neutral synonymous language.*

Given the important and widespread policy implications and ongoing applications of such an intervention, we investigate if, why, and for whom it increases application behavior. Gaucher et al. argue that "…masculine wording likely signals that there are many men in the field and alerts women to the possibility that they do not belong" (p. 111). That is, prior literature focuses mainly on how masculine language triggers gender category misfit among women (5, 18). Thus, gender debiasing interventions could eliminate women's feelings of misfit when confronted with masculine language; consequently, women's anticipated belonging and intention to apply should increase. We extend these predictions even further by incorporating continuous variation within gender categories.

*For instance, some words may not appear explicitly masculine but may still carry indirect masculine connotations. Furthermore, what is considered "masculine" or "feminine" language is contextual: it may shift over time, vary across cultures, and depend on the measurement method used. We discuss these limitations in our General Discussion.

Author affiliations: [a]Anderson School of Management, University of California, Los Angeles, CA 90095; and [b]Department of Management, University of Toronto, Mississauga, ON L5L 1C6, Canada

Author contributions: J.C.H. and S.K.K. designed research; J.C.H. and S.K.K. performed research; J.C.H. analyzed data; and J.C.H. and S.K.K. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: Joyce.he@anderson.ucla.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2409854122/-/DCSupplemental.

Published February 10, 2025.

We consider two dimensions of variability within gender-category—prototypicality and identification—to theorize how not only women, but also certain men, experience incongruence with masculine language and could benefit from a gender debiasing intervention. Individuals who are more prototypical of a category experience stronger effects of categorization (e.g., 19, 20). For instance, male-typical men experience fewer belonging concerns in majority-male contexts (21–24). Related, gender identification describes how central one's gender category is to one's self-identity (25–28). Individuals who are strongly identified are more psychologically attached to their gender category (28, 29) and exhibit gender-norm congruent behavior and self-stereotyping that further reinforce identification (25, 29).

Considering variation both between and within gender categories, we posit that men who are less gender-identified or less gender-typical also feel misaligned with the strong male identity espoused by masculine default language (24, 30). Therefore, we predict that a gender debiasing intervention that removes masculine language from job postings and replaces it with gender-neutral language will increase application rates among women and among weakly identified or less male-typical men. Moreover, we theorize that the predicted increase in application rates stems from increases in anticipated belonging (31–35). Overall, these predictions describe the broader effects of a debiasing intervention on all individuals who do not fit with the masculine default, including both women and men.[†]

We test our predictions across four studies, employing quantitative and qualitative methods in lab and field settings. We begin with a pilot study that inductively and qualitatively explores our core assumption that male gender role congruity varies along a continuum of gender identification. We then test our predictions in the field, first examining the relationship between masculine language in job advertisements and gender diversity of applicant pools (Study 1), and in turn using these insights to design a quasi-experimental test of our "debiasing" intervention using a real job advertisement for a historically male-dominated position (Study 2). Finally, we replicate these findings in a controlled online field experiment and identify anticipated belonging as the mechanism driving these effects (Study 3). Our findings overall demonstrate that replacing masculine language with gender-neutral language benefits people whose gender identification and gender-typicality do not fit within a masculine "blueprint."

Analysis code, deidentified data, and research materials are available on OSF (36),[‡] as is our *SI Appendix*. Some materials (i.e., job postings used in Studies 2 and 3) are redacted for privacy reasons. Pilot study qualitative analyses were not preregistered because themes were coded inductively, but quantitative analyses were preregistered at https://aspredicted.org/F26_HVS. Studies 1 and 2 were not preregistered because the data were secondary data; Study 3 was preregistered and is available at https://aspredicted.org/UPW_ZJM. This study was approved by the research ethics board at the University of Toronto (protocol ID 35345). Studies 1 and 2 were secondary data so we did not obtain consent directly from participants. In the pilot study and Study 3, participants provided informed consent and were debriefed.

Pilot Study

In our pilot study (n = 793), we explored men and women's inferences about and reactions to masculine language in job postings. We provided examples of masculine language in job advertisements and asked participants their reactions in open-ended questionnaires. We found that while women overall reported being less likely to apply to job postings with masculine language, almost one third of men also had a negative response to the language; this negative response was pronounced among men who self-reported as being weakly identified with their gender. Two main themes emerged for how masculine language discourages men: 1) misalignment with the masculine attributes of the job [e.g., "I would probably not (apply), I don't find myself to fit with expected masculine traits"] and 2) misalignment with the masculine culture (e.g., "In my experience men have stricter requirements for how masculine other men should be so I'd worry about not getting along with coworkers"). These findings provide initial support for our core assumption that in addition to women, certain men may also experience incongruence with masculine language in job postings that stem from belonging concerns with masculinity.

In this study, we also empirically validated a name-based predictor of gender category and gender-typicality (gender package in R; 37, 38), which we use in Studies 1 and 2, where self-reported gender information is unavailable. Specifically, given a first name as input, this algorithm searches historical databases to generate a "proportion male" estimate—a continuous variable ranging from 0 to 1 that represents the proportion of people with that name who self-identified as men.[§] Based on this estimate, the algorithm assigns a binary gender classification (greater than 0.50 proportion male = "man"; less than 0.50 proportion male = "woman"; or "unknown" if the name does not appear in the database).

Theoretically, the gendering of first names may reflect individuals' associations with their gender category. Although individuals do not typically choose their own name, the gendering of a name can shape self-perceptions of gender typicality (and in turn, gender identification) through socialization processes and self-fulfilling prophecy mechanisms (39). This can occur, for example, through gender-consistent or inconsistent treatment from others based on apparent gender typicality, which over time inform individuals' identification with their gender category (39–41). Moreover, parents who choose highly gender-typical names for their children may also shape their children's identification and behavior in line with such gender norms and expectations (42, 43). Indeed, research finds that men with less male-typical first names report less male-typical childhood behavior and fewer masculine personality traits (44). We elaborate on this theoretical connection in the *SI Appendix*.

Accordingly, from the gender package, we use gender category predictions to infer gender category, and the continuous "proportion male" estimate to infer perceptions of gender typicality[¶] based on names, which we term "inferred name-based male-typicality score." To validate this tool, in the same survey, participants indicated the name that they use on job applications and answered some gender-based questionnaires. We found that the gender package accurately predicted self-reported gender category 90.8% of the time, and the algorithm-generated inferred name-based male-typicality score was very highly correlated with participants' self-assessments of name-based male typicality (r = 0.87, P < 0.001), which was in turn predictive of men's gender identification (although there was no direct effect of inferred name-based male-typicality score on gender identification). We

[†]We acknowledge that our arguments are nonetheless framed according to binary gender categorization (16). We discuss extensions of our work beyond this binary framework in our discussion section.

[‡]Due to space constraints, some of these methodological details are also contained in our *SI Appendix* on OSF: https://osf.io/mub7n.

[§]A name with a proportion male score of 0 indicates that no individuals with that given name self-identified as a man, whereas a name with a score of 1 indicates that all individuals with that given name self-identified as a man.

[¶]From a frequentist and descriptive norm-based perspective, individuals may perceive names as more or less male-typical based on the observed proportion of women to men who share that name. These observations contribute to broader culturally perceptions of certain names as more or less male-typical, which in turn can shape and become internalized in individuals' self-perceptions of gender-typicality.

interpret these results as evidence supporting the validity of the inferred name-based male-typicality score for predicting self-assessed male-typicality based on names, though not necessarily gender identification.

Study 1

Having explored our core assumptions and validated our name-based predictor of gender category and typicality in the pilot study, in Study 1, we examine whether natural variation in gendered language in job advertisements predicts variation in gender composition of the corresponding applicant pool.

Method.

Participants and data. We obtained a database of 576 unique job advertisements and matching applicant data from 32,834 applicants from a large public company in Canada, so we therefore could not perform an a priori power analysis. The jobs were posted over 16 mo from September 2015 to January 2017.

Measures.

Gendered language. We used Linguistic Inquiry Word Count (LIWC; 45), to analyze the job advertisements for gendered language using an established dictionary (9). These gendered language dictionaries composed of agentic and communal word (e.g., individualistic, competitive; committed, supportive), and masculine and feminine trait words (e.g., ambitious, assertive; compassionate, understanding) (46–49).

Gender composition of applicant pool. We used applicant's self-reported gender where available (just over half of the sample, 54.4%). For applicants who did not disclose their gender, we used the gender package in R (v0.5.3; 38) that we validated in the pilot study to predict their gender from their first names. The total applicant pool consisted of 13,086 (39.8%) women, 17,725 (54.0%) men, and 2,023 (6.2%) applicants of unknown gender. To examine gender composition of the applicant pool for each job, we divided the number of women who applied by the total number of men and women who applied ($M_{\text{percentage female applicants}}$ = 38%, SD = 22%).[#]

Name-based male-typicality of applicant pool. We operationalized name-based male typicality using the inferred name-based male-typicality score (i.e., "proportion male" variable) generated from the gender package (method validated in pilot study). The overall distribution of this variable was bimodal between names that were either 0% (coded as women) or 100% male-typical (coded as men), with fewer names falling in between (1 to 99%). Because the bimodal distribution confounds inferred name-based male-typicality score with gender category,[||] we further differentiated names based on name type: comparing names categorized as 0% or 100% male-typical against names that ranged from 1% to 99%. We focus on the test of our predictions along a continuum of fit primarily within these continuous name types (n = 4,724) for a conservative test of our effects.

Job level. This organization groups jobs by level in terms of similar compensation and qualifications, from 1 (entry-level) to 9 (senior management). The jobs in our database ranged from level 2 to level 9 (M = 5.62, SD = 1.79).

Job context. There were 12 different departments in our data (e.g., Information Technology, Human Resources). We used the proportion of women to men that applied to each department to determine its gender-type (male-dominated: female applicants <45%; gender-neutral: female applicants 45 to 55%; female-dominated: female applicants >55%). All departments in the dataset were male-dominated (M = 37.1% women, SE = 0.95%; range = 16.9 to 40.0%), except Human Resources/Marketing (54.5% women, SE = 2.4%), $F(1, 571)$ = 10.74, P = 0.001. Thus, job gender-type was coded as 0 = male-dominated and 1 = non-male-dominated.

Controls. We control for word count and words per sentence to account for overall length differences across job advertisements. We also control for job level and feminine language.

Analysis plan. We used hierarchical linear modeling (HLM). For hypotheses at the job posting level, we include proportion of women in the applicant pool per job posting as our Level 1 outcome variable, masculine language in the job posting as the Level 1 independent variable, with a random intercept for job level (Level 2 variable; ICC = 0.22). For hypotheses at the male-typicality level, we include male-typicality of names as our Level 1 outcome variable and masculine language as our Level 2 predictor (ICC = 0.09), with a random intercept for job posting (Level 2).

Results. Descriptive statistics and zero-order correlations are available in the *SI Appendix*. The proportion of masculine language in job advertisements is negatively related to the proportion of women in the applicant pool (b = −0.034, SE = 0.01, P = 0.003) (Model 1; Table 1*A*). At one SD below the mean of masculine language, women comprise 40.0% of the applicant pool (SE = 3.6%), whereas women comprise 34.1% of the applicant pool at one SD above the mean (SE = 3.6%). Adding job gender-type as an additional predictor (Model 2) does not reduce the effect of masculine language, and treating it as a moderator (Model 3) demonstrates that masculine language was negatively related to the proportion of women in the applicant pool for male-dominated jobs (b = −0.045, SE = 0.01, P < 0.001; Fig. 1, Panel *A*), but the relationship was nonsignificant for other jobs (b = 0.016, SE = 0.035, P = 0.64). These effects hold with additional controls (Model 4; Table 1*A*), and when we only examine the subset of applicants who self-identified their gender (*SI Appendix*).

Next, we examined the relationship between masculine language in the job posting and applicants' inferred name-based male-typicality score (Model 1, Table 1*B*), focusing particularly among applicant names between 1% to 99% on inferred name-based male-typicality (i.e., with more continuous variation on male-typicality). For men, the higher their inferred name-based male-typicality score, the more likely they were to apply to job postings with higher proportions of masculine language (b = 0.010, SE = 0.004, P = 0.004). These results held with controls (Model 2, Table 1*B*) and remained robust when we ran the same analyses with only applicants for male-dominated jobs. The inferred name-based male-typicality score for women also positively predicted masculine language in the postings for which they applied (b = 0.018, SE = 0.005, P < 0.001) (Model 3 to 4, Table 1*B*). Overall, these results suggest that name-based male typicality (reflecting within-gender-category variation) mattered for both men and women (Fig. 1, Panel *B*).

Study 2

We partnered with a large Canadian investment company to administer our intervention on one historically male-dominated entry-level "feeder" position that was reposted roughly every month, allowing us to compare the applicant pool before and after our intervention.

Method.

Participants. The sample included monthly applicant data from September 2015 to October 2019. Our intervention (the debiased

[#]Results hold when we examine the number of female applicants divided by the total number of applicants (including unknown gender) and when we examine proportion of male applicants; results available in our *SI Appendix*.

[||]19,898 (80.8%) of the names in our dataset were either 0% male-typical (thus categorized as a woman) or 100% male-typical (thus categorized as a man). We show detailed distribution and examples of names in our our *SI Appendix*.

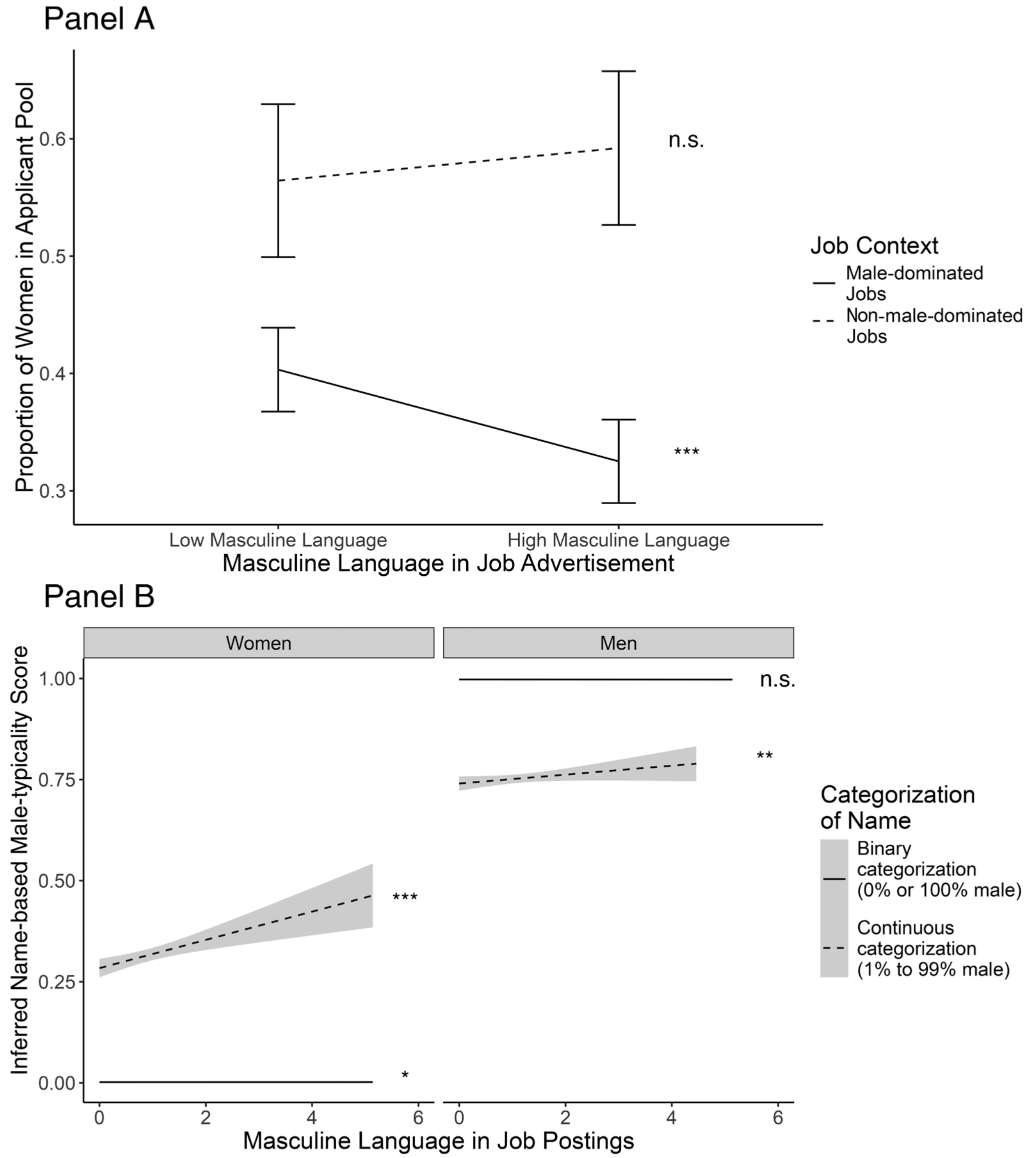

Fig. 1. Study 1 results of masculine language in job ads predicting applicant pools. *Note.* Panel *A* displays the Study 1 cross-level moderation of job gender-type on the effect of masculine language in job postings on the proportion of female vs. male applicants. Panel *B* illustrates the relationship between inferred name-based male-typicality score of an applicant's name and the proportion of masculine language in the job postings they applied for, according to whether the name was categorized as 0% or 100% male-typical (binary categorization), or 1 to 99% male-typical (continuous categorization).

job posting) was administered in January 2019. The organization does not ask applicants for their gender, so we predicted it using first names (i.e., using the gender package in R as in Study 1). Of 3,503 total applicants, 2,288 were coded as men (65.3%), 675 were coded as women (19.3%), and 540 were not categorizable and therefore unknown (15.4%).

Materials and procedure. To develop the intervention, we first identified stereotypically masculine words in the job posting using LIWC and the same gendered language dictionary as in Study 1. We additionally coded "entrepreneurial" as masculine due to its strong stereotypical association with men (50), and the phrases "drive results" and "passion for capital markets" as masculine due to their similarity to masculine terms in the dictionary (*SI Appendix* for details on material development). The baseline job posting had 1.52% masculine language and 0.61% feminine language. We replaced these masculine words with synonymous gender-neutral words (e.g., motivation to achieve results, interest in capital markets), working closely with the organization's HR team to verify that the underlying substantive content remained unchanged. The debiased job posting had 0.6% masculine language and 0.89% feminine language. A validation experiment ($n = 390$) verified that the debiased job posting significantly reduced perceptions of masculinity compared to the original job posting without increasing perceptions of femininity. There were no other changes made to the job posting. The redacted job postings and validation details are available on our OSF site.

Measures.

Gender composition of the applicant pool. As in Study 1, we compute percentage of women in the applicant pool by dividing the number of female applicants by the total number of male and female applicants for each month (M = 22%, SD = 6%). Analyses using alternative gender composition measures provide similar results and can be found in our *SI Appendix*.

Name-based male-typicality of applicants. We used the same inferred name-based male-typicality score from Study 1.

Analysis. For applicant-level analyses, i.e., to test hypotheses about name-based male-typicality, we examined changes in overall inferred name-based male-typicality scores pre- and postintervention by name-type and applicant gender, as in Study

Table 1. Hierarchical linear model results for study 1

DV: Proportion of women in applicant pool				
Variable	Model 1	Model 2	Model 3	Model 4
Masculine language	−0.034 (0.01)**	−0.040 (0.01)***	−0.045 (0.01)***	−0.046 (0.01)***
Job gender-type (non-male-dominated)		0.215 (0.05)***	0.214 (0.05)***	0.211 (0.05)
Masculine language * Job gender-type			0.061 (0.04)+	0.067 (0.04)+
Feminine language				−0.014 (0.02)
Word count				−0.0001 (0.00003)**
Words per sentence				−0.0004 (0.0007)
Job level				−0.035 (0.006)***
Constant	0.37 (0.03)***	0.36 (0.04)***	0.36 (0.03)***	0.56 (0.03)***
N	573	573	573	573
Pseudo R^2	0.009	0.050	0.052	0.057
Δ Pseudo R^2	0.009	0.040	0.002	0.005
Model comparison	χ2(1) = 9.02***	χ2(1) = 24.52***	χ2(1) = 2.91+	χ2(4) = 27.1***

DV: Inferred name-based male-typicality score				
	Male applicants		Female applicants	
Variable	Model 1	Model 2	Model 3	Model 4
Masculine language	0.0002 (.002)	−0.0005 (0.002)	0.007 (0.003)*	0.008 (0.003)**
Name type [binary (0) vs. continuous (1)]	−0.199 (0.006)***	−0.200(0.006)***	0.205 (0.006)***	0.204 (0.006)***
Masculine language * Name type	0.010 (0.004)**	0.010(0.004)**	0.011 (0.005)*	0.011 (0.005)*
Feminine language		−0.003 (0.003)		−0.007 (0.004)
Word count		0.0000 (0.0000)		−0.0000 (0.0000)
Words per sentence		0.0001 (0.0001)		−0.0005 (0.0002)*
Constant	0.981 (0.003)***	0.978 (0.006)***	0.018 (0.003)***	0.038 (0.008)***
N	14,200	14,200	10,422	10,422
Pseudo R^2	0.1864	0.1863	0.1810	0.1812
Δ Pseudo R^2	0.1864	−0.0001	0.1810	0.0002
Model comparison	χ2(3) = 2917***	χ2(3) = 2.674	χ2(3) = 2110***	χ2(3) = 8.51*

Note. + indicates $P < 0.10$; * indicates $P < 0.05$; ** indicates $P < 0.01$; *** indicates $P < 0.001$. The table displays HLM regression estimates with SE in parentheses. We include a pseudo R^2 as calculated by the change in within-group variance when additional predictors are added (Bliese, 2002). Job Gender-Type is coded as 0 = male-dominated, 1 = non-male-dominated.

1. For job-posting level analyses, i.e., to test hypotheses about the gender composition of the applicant pool, we used an interrupted time-series analysis (via a segmented regression analysis) to examine applicant pool changes as a function of the job description intervention with the following model specification:

$$Proportionwomen_t = b_0 + b_1\,time + b_2\,intervention + b_3\,timesince + e_t.$$

This model estimates an intercept (b_0), the slope of proportion of women over time before the intervention (b_1), the immediate change in intercept before and after the intervention (b_2), the slope of proportion of women over time after the intervention (b_3), and a residual error term (e_t). Segmented regression examines the immediate change in level and slope of a variable after an event, while accounting for the data's temporal nature. This approach is commonly used in other fields and has recently been introduced to psychology and organizational behavior (51, 52).

In total, we had data from 42 timepoints: 32 dates (months) and 2411 applicants before the intervention; 10 dates (months) and 1092 applicants after the intervention. To correct for any autocorrelation in time series data, we use the Newey–West estimator (53); we also examine the heteroskedasticity and autocorrelation consistent (HAC) estimation of the covariance matrix using the Sandwich function in R (54, 55) as a robustness check. We report the main results using the Newey–West estimator but only discuss findings significant across both estimators.

Results. First, we examined the overall effects of the debiasing intervention. There was an immediate increase in the proportion of female applicants after the intervention, $b = 0.05$, $SE = 0.01$, $P = 0.001$, with a nonsignificant decrease over time postintervention, $b = -0.001$, $SE = .0008$, $P = .09$ (Fig. 2, Panel *A*). These results suggest that the intervention garnered an immediate increase in the proportion of female applicants by approximately 4% (7 more women each month), an effect that was sustained 10 mo after the intervention. To better understand what was driving this proportional change, we next analyzed the total number of female and male applicants separately (details in *SI Appendix*). There was a robust and sustained increase in the total number of women and men applying to the job after the intervention that did not return to baseline (Fig. 2, Panel *B* and *C* respectively), corresponding to an increase in the total number of applicants each month.

At the applicant-name level, the intervention was associated with a significant decrease in the inferred name-based male-typicality score of male applicants with names that ranged from 1 to 99% male typical ($b = -0.012$, $SE = 0.006$, $P = 0.04$). For women, the intervention also decreased inferred name-based male-typicality score among continuously categorized (1 to 99% male typicality) female names ($b = -0.078$, $SE = 0.010$, $P < 0.001$). The debiasing intervention was associated with an overall increase

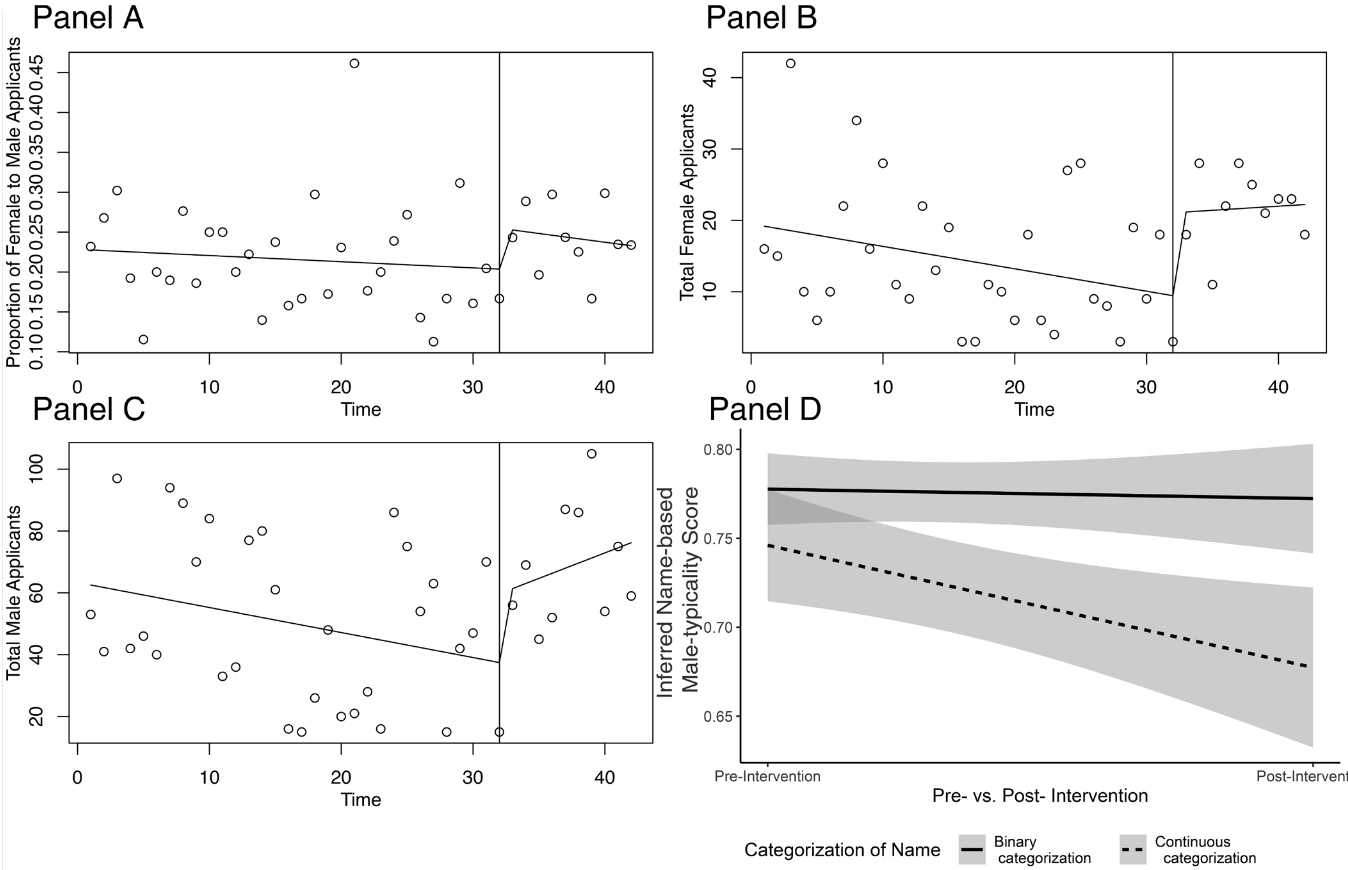

Fig. 2. Study 2 results of the debiasing intervention on applicant outcomes. *Note.* Panels *A*–*C* show the effect of the debiased job advertisement over time on the proportion of female to male job applicants (Panel *A*), total female applicants (Panel *B*), and total male applicants (Panel *C*). The horizontal axis represents time points (each time point is approximately 1 mo), and the vertical lines show when the intervention was applied (at time point = 32). Panel *D* displays the comparison of the average inferred name-based male-typicality score of applicant names before and after the intervention.

in all applicants with a lower inferred name-based male-typicality score, as shown in Fig. 2, Panel *D*.

Study 3: Experimental Test with Mechanisms

Method.

Participants. An a priori power analysis indicated that we would need 790 participants to detect a small effect size with 80% power. We recruited 826 participants in the United States on Prolific; as preregistered, we excluded participants who failed the comprehension check (*n* = 36). Our final sample was 790 participants; 376 (47.6%) self-identified as men, 396 (50.1%) self-identified as women, and 18 (2.3%) chose another option. On average, participants were 33.28 years old (*SD* = 11.52) and were primarily white (*n* = 521; 65.9%).[**]

Materials and procedure. Participants were randomly assigned to view either the original or debiased job posting from Study 2 (modified to include a fictional organization name). After seeing the job posting, participants completed the measures below and a brief demographic questionnaire.

Measures.

Manipulation check. Participants indicated their perceptions about the masculinity and femininity of the organization's ideal applicant. Participants in the debiased job posting condition perceived the ideal candidate as less masculine (*M* = 5.15, *SE* = 0.037) than in the control condition (*M* = 5.25, *SE* = 0.037), $F(1, 787) = 3.90$, $P = 0.049$, $\eta_p^2 = 0.005$); perceptions of the ideal candidate's femininity did not differ across conditions, $F(1, 787) = 0.77$, $P = 0.38$, $\eta_p^2 = 0.001$). Moreover, perceived representation of women relative to men in the position did not differ between conditions ($P = 0.64$), supporting our theorizing about masculine language signaling a more continuous notion of masculinity.

Dependent variables. Job appeal was measured using a previously validated scale (1 = weak appeal; 7 = strong appeal) (9; α = 0.91; items in *SI Appendix*). Likelihood of applying was measured by the item "How likely are you to apply for a position like this one?" on a scale from 1 = very unlikely to 7 = very likely. For brevity, we focus on job appeal in our reporting of the results, noting whether results for application intent were consistent or different.

Anticipated belonging. Anticipated belonging was measured with a 4-item scale: "I could fit in well in this position," "I am similar to the people who work in this position," "My values and this company's values are similar," "The type of people who would apply for this position are very different from me" (reverse-coded) (9; α = 0.88).

Gender identification. Gender identification was measured using a 4-item measure, e.g., "Being a woman/man is an important reflection of who I am" (27; α = 0.93).

Job context. To closely replicate the field setting of Study 2 (i.e., applicants applying for a position within a male-dominated industry, in the case of Study 2, finance), we focus on applicants who currently work or have previously worked in male-dominated industries. We focus on these participants because gender-relevant cognitions are elicited by specific contexts and cues (56)—male-

[**]Additional ethnicities: 11.4% East and Southeast Asian (n = 90), 4.4% South Asian (n = 35), 0.6% Middle Eastern (n = 5), 6.2% Black/African American (n = 49), 10.2% Hispanic/Latino (n = 81), 2.0% Black/Caribbean (n = 16), 4.6% Asian/Pacific Islander (n = 36), 1.3% Native American (n = 10), 1.3% Other (n = 10). Total proportions exceed 100% because participants could select more than one option.

identity belonging concerns are exacerbated among men and women embedded in male-dominated job contexts because they are more vigilant for and sensitive to cues that signal lower gender-based belonging. We used the Bureau of Labor Statistics' major occupational categories (https://www.bls.gov/cps/cpsaat11.htm) to code participants' current occupations as female-dominated (>55% female), male-dominated (<45% female), or gender-neutral (45 to 55% female). Overall, 253 (32.9%) of participants worked in a female-dominated industry, 228 (29.6%) in a male-dominated industry, and 288 (37.4%) in a neutral-typed industry. Treating this variable as continuous does not change our findings.

Control variables. We controlled for perceived ability as an alternative mechanism predicting job pursuit (9) in our indirect effect analyses.

Analysis. Our preregistration specified testing interactions between gender and condition on job appeal and likelihood of applying; while this two-way interaction was not significant for our dependent variables (*P*s > 0.17), we also preregistered additional moderations by gender identification and participant industry, which became the focus of our analysis. For certain hypotheses, we tested interactions between gender, condition, industry-type, and gender identification. Given the complexity of understanding 4-way interactions, we deviated from the preregistration to create a subset of participants from male-dominated industries and examined the three-way interaction between gender, gender identification, and experimental condition within this data subset, with a caveat that we are relatively underpowered for such a complex design. To streamline the reporting and interpretation of results, we focus on this subset of participants who reported working in a male-dominated industry in the results section. The analyses with the full sample, including with participant industry as an additional moderator for the full four-way interaction, are reported in our *SI Appendix*. Notably, our hypothesized interactions and simple contrasts remain significant. Table 2 displays our regression results (simple, indirect, direct, and total effects) and Fig. 3 displays our results by gender category (Panels *A* and *C*), then broken down by gender identification (Panels *B* and *D*) among participants working in a male-dominated industry. Condition was coded as 1 = debiased condition, 0 = control condition, and we use the Bonferroni adjustment to account for multiple comparisons.

Results. There was a significant interaction between participant gender and job posting condition (b = 0.77, SE = 0.37, P = 0.037 for job appeal; b = 1.21, SE = 0.52, P = 0.021 for likelihood of applying). As predicted, women currently working in male-dominated industries found the job posting more appealing in the debiased condition (M = 4.29, SE = 0.20) than in the control condition (M = 3.36, SE = 0.21; b = 0.92, SE = 0.29, P = 0.01) (Fig. 3, Panel *A*). Overall, the gap between men and women in the control condition (b = −1.28, SE = 0.26, P < 0.001) was reduced and no longer significant in the debiased condition (b = −0.52, SE = 0.26, P = 0.27). The debiasing intervention had a significant and positive indirect effect on job appeal through belonging (Table 2). All results held with intent to apply as the outcome variable.

Next, the three-way interaction between gender identification, gender, and condition was also significant (b = −0.87, SE = 0.29, P = 0.003 for job appeal; b = −1.04, SE = 0.42, P = 0.014 for likelihood of applying). Overall, perceived job appeal did not differ between men (categorically) in the debiased (M = 4.84, SE = 0.15) and control (M = 4.66, SE = 0.15) conditions (Fig. 3, Panel *A*) (P = 0.99). However, probing by gender identification revealed that men at 1 SD below the mean of gender identification saw the debiased job posting as marginally significantly more appealing (P = 0.053) than the control; although they were descriptively more likely to apply to the debiased job, this difference was no longer significant after correcting for multiple comparisons (P = 0.28) (Fig. 3, Panel *C* and *D*). For these men, there was a significant and positive indirect effect of debiased job posting condition on job appeal through increased belonging (Table 2; results hold for intent to apply as the outcome variable). Conversely, for men at 1 SD above the mean of identification, there was no significant difference in perceived job appeal nor application interest in terms of job posting condition (*P*s > 0.20). Examining gender identification among women, we observed no difference in job appeal (P = 0.99) and likelihood to apply (P = 0.99) between the debiased and control postings among weakly identified women. Conversely, strongly identified women found the debiased posting more appealing (b = 1.32, SE = 0.39, P = 0.005) and were marginally more likely to apply to it (b = 1.38, SE = 0.55, P = 0.08) than the control posting; the indirect effect through anticipated belonging was also significant (Table 2).

General Discussion

We identify and test a subtle, light-touch intervention that replaces masculine language with gender-neutral (rather than feminine) synonyms, demonstrating modest but robust effects. These findings not only demonstrate support for policy implications of Gaucher et al.'s work (9) in the field, but also extend it in important ways. We demonstrate broader effects of the intervention on women and men who do not fit with the blueprint of masculinity. This underscores the importance of conceptualizing incongruence along a continuum of gender identification rather than focusing only on discrete gender categories to more comprehensively understand the effect of gender debiasing interventions. Whereas gender inequality has been understood primarily as the categorical sorting of men into male-dominated jobs, our findings reinforce recent theoretical advances that suggest a broader reconceptualization of gender segregation as the sorting of masculine-identifying and more male-typical individuals of all genders into masculine jobs (4). We contribute to this literature on masculine defaults by demonstrating one way that cultural cycles of gender bias are perpetuated and how they can be disrupted: System-level policy changes can affect the composition and gender association of individuals who approach and enter male-dominated spaces, which could ultimately shift masculine default cultures.

Practically, we propose a concrete intervention to increase gender diversity, and provide a comprehensive test of mechanisms, effect sizes, and boundary effects in lab and field settings. For instance, our findings suggest that the intervention led to increased application rates among women and men who are weakly identified with masculinity or who appear less male-typical in male-dominated job contexts (i.e., for male-dominated jobs and among applicants who currently work in male-dominated industries; *SI Appendix* for an overview of job context operationalizations across studies). Given recent work documenting no effect of gendered language on applicant behavior in gender-neutral or female-typed jobs (57), using the framework of masculine defaults to conceptualize masculine language as signals of masculinity in already male-dominated domains may most aptly capture its effects on applicant behavior. Our results suggest that gender debiasing interventions in these male-dominated domains could effectively increase gender diversity among applicants, with the caveat that we observe relatively small effect sizes. Given the cost–benefit trade-off, this remains an easy and cost-effective way for organizations to instigate changes to their recruitment process. Importantly, we show that gender-inclusive workplace cultures

Table 2. Moderated mediation models for Study 3

Men in male-dominated industries

Outcome variable	*Mediator*	*a path*	*b path*	*Indirect effect*	*Direct effect*	*Total effect*
Job appeal	Belonging	0.23 (0.23)	0.70 (0.06)***	b = 0.16, SE = 0.17, P = 0.35	b = −0.02, SE = 0.12, P = 0.87	b = 0.15, SE = 0.22, P = 0.48 (P = 0.99)
	Ability	0.15 (0.27)	0.11 (0.05)*	b = 0.02, SE = 0.03, P = 0.63		
Likelihood of applying	Belonging	0.23 (0.23)	0.75 (0.09)***	b = 0.17, SE = 0.18, P = 0.34	b = −0.32, SE = 0.20, P = 0.11	b = −0.10, SE = 0.31, P = 0.74 (P = 0.99)
	Ability	0.15 (0.27)	0.31 (0.08)***	b = 0.05, SE = 0.09, P = 0.60		

Women in male-dominated industries

Outcome variable	*Mediator*	*a path*	*b path*	*Indirect effect*	*Direct effect*	*Total effect*
Job appeal	Belonging	0.86 (0.31)***	0.70 (0.06)***	b = 0.60, SE = 0.21, P = 0.004	b = 0.24, SE = 0.17, P = 0.15	b = 0.92, SE = 0.29, P = 0.002 (P = 0.01)
	Ability	0.77 (0.35)*	0.11 (0.05)*	b = 0.08, SE = 0.06, P = 0.15		
Likelihood of applying	Belonging	0.86 (0.31)***	0.75 (0.09)***	b = 0.65, SE = 0.24, P = 0.006	b = 0.22, SE = 0.27, P = 0.42	b = 1.10, SE = 0.41, P = 0.008 (P = 0.049)
	Ability	0.77 (0.35)*	0.31 (0.08)***	b = 0.24, SE = 0.12, P = 0.06		

Men x gender identification in male-dominated industries through belonging

Outcome variable	*Identification*	*a path*	*b path*	*Indirect effect*	*Direct effect*	*Total effect*
Job appeal	High	−0.24 (0.31)	0.67 (0.06)***	b = −0.16, SE = 0.20, P = 0.43	b = −0.16, SE = 0.18, P = 0.38	b = −0.36, SE = 0.30, P = 0.22 (P = 0.99)
	Low	0.78 (0.29)**	0.67 (0.06)***	b = 0.52, SE = 0.22, P = 0.02	b = 0.12, SE = 0.17, P = 0.46	b = 0.74, SE = 0.28, P = 0.008 (P = 0.05)
Likelihood of applying	High	−0.24 (0.31)	0.72 (0.09)***	b = −0.17, SE = 0.22, P = 0.43	b = −0.61, SE = 0.29, P = 0.04	b = −0.91, SE = 0.42, P = 0.03 (P = 0.20)
	Low	0.78 (0.29)**	0.72 (0.09)***	b = 0.56, SE = 0.23, P = 0.02	b = −0.02, SE = 0.27, P = 0.95	b = 0.79, SE = 0.39, P = 0.05 (P = 0.28)

Women x gender identification in male-dominated industries through belonging

Outcome variable	*Identification*	*a path*	*b path*	*Indirect effect*	*Direct effect*	*Total effect*
Job appeal	High	1.01 (0.41)*	0.67 (0.06)***	b = 0.68, SE = 0.30, P = 0.02	b = 0.56, SE = 0.23, P = .02	b = 1.32, SE = 0.39, P < .001 (P = 0.005)
	Low	0.35 (0.50)	0.67 (0.06)***	b = 0.24, SE = 0.29, P = 0.42	b = −0.17, SE = 0.28, P = 0.56	b = 0.12, SE = 0.48, P = 0.80 (P = 0.99)
Likelihood of applying	High	1.01 (0.41)	0.72 (0.09)***	b = 0.73, SE = 0.34, P = 0.03	b = 0.43, SE = 0.38, P = 0.26	b = 1.38, SE = 0.55, P = 0.01 (P = 0.08)
	Low	0.35 (0.50)	0.72 (0.09)***	b = 0.25, SE = 0.31, P = 0.42	b = −0.04, SE = 0.46, P = 0.93	b = 0.35, SE = 0.68, P = 0.61 (P = 0.99)

Note. Moderated mediation models for job condition on job appeal and likelihood of applying through anticipated belonging and ability fit for women and men in male-dominated industries, further separated by low and high gender identification in Study 3. For the moderated mediations further separated by gender identification, the simultaneous indirect effects through ability (in addition to belonging) are not reported in the table for brevity. For total effects, Bonferroni-adjusted P values are presented in parentheses.

benefit a broad group of individuals, including men, which may increase interest and action among men and, ultimately, garner broader support for gender diversity initiatives.

Limitations and Future Directions. Although our findings suggest that the intervention may be particularly effective in male-dominated job contexts, we use different operationalizations of job context across studies (i.e., the occupational field of the job versus the applicant's current industry) due to limited data access. As we elaborate in our *SI Appendix*, we posit that both operationalizations capture how context amplifies the salience of identity-related concerns, making applicants in such contexts especially vigilant and sensitive to cues that signal their lack of belonging (i.e., via masculine language). However, we cannot make causal claims about job context since we do not provide a causal test of this moderator. Future work can extend these findings by experimentally manipulating job context (e.g., by comparing the intervention in male-dominated vs. female-dominated fields), to test our hypotheses causally, ideally using larger field samples.

Our predictions focused on gender typicality and identification as important within-category predictors for men, but our results show that these processes also moderate women's responses to masculine language. Women who are more female-typical and gender-identified were especially likely to respond negatively to masculine language, and thus exhibit even larger increases in belonging and application behavior in response to a gender debiasing intervention. Future research can continue to examine how gender identification and gender categorization interact to impact women's and men's sense of belonging.

A major limitation of our work is the focus on people who identify as women and men. Gender self-categorization and identification extend beyond this binary, and more research is needed to examine how people of all genders react to gender cues in job

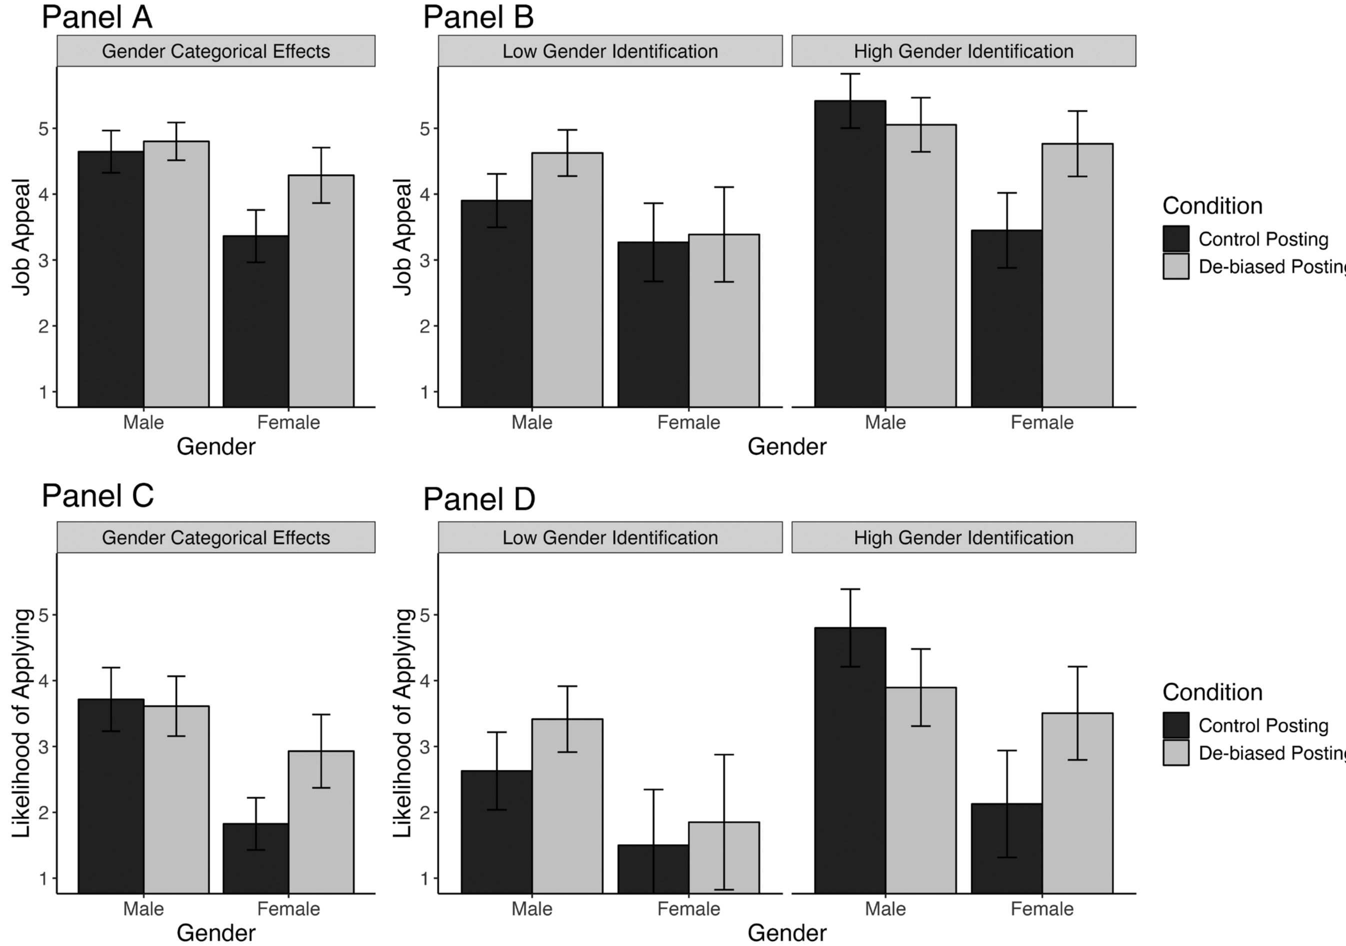

Fig. 3. Study 3 experimental results by gender and identification for participants working in male-dominated industries. *Note.* Study 3 participants' ratings of job appeal (Panel *A* and *B*) and ratings of likelihood to applying to the job (Panel *C* and *D*) among subset of participants who reported working in a male-dominated industry. Panel *A* and *C* display the results for participant categorical gender (i.e., man vs. woman, collapsing across gender identification) by job posting condition for job appeal and likelihood of applying to the job respectively. Panel *B* and *D* display the results for participant categorical gender by job posting condition further moderated by self-reported gender identification (i.e., 1 SD below the mean of gender identification and 1 SD above the mean of gender identification) for job appeal and likelihood of applying to the job respectively. Error bars represent 95% CI.

advertisements. Of note, qualitative responses from individuals who identified as nonbinary in our pilot study suggest that our debiasing intervention may signal more inclusion to "…anyone who is considered less masculine—whether it be because the person is a woman, is trans, [or] is simply a more feminine man." One participant who identified as nonbinary said: "I definitely would steer away from any job postings containing more masculine language because I'd like to work for a place that didn't care about genders." Future research should explore how the current effects relate to individuals who identify with gender categories beyond men and women or who have more flexible gender identities.

Our findings are also limited by our use of the gender package in R to probabilistically infer gender, which is less reliable than self-report and relies on name databases from the United States, limiting generalizability across locations and cultures. Moreover, while we theorize that our algorithm-based inferred name-based male-typicality score may capture perceptions of male-typicality, future research should seek field data containing applicants' self-reported gender and self-assessed within-gender variability to verify our results. We also rely on a theory-driven dictionary approach to measuring gendered language, which may not comprehensively capture its nuance. Future research can improve these methods by leveraging advances in natural language processing and AI to better assess gendered language using bottom–up approaches that account for nuance (e.g., (57, 58), considering variations across culture, time, and the strength of gender association (59). Additionally, future research should test whether our theoretical tenets about gender debiasing apply to other measures of masculine language across different contexts. This is an important area for future research, especially as gender is increasingly understood as fluid and continuous, extending beyond fixed binary categories and psychological processes within this binary.

Related, our intervention may also change the applicant pool along other dimensions, such as race, socioeconomic status, personality, competitiveness, etc. However, we were unable to provide a more granular test of the effect of debiased job postings beyond gender with the current data. Future work can explore the effect of debiasing across different social groups and at their intersections.

Finally, we delineate key limitations of our intervention approach and their implications for organizations considering its implementation. Given the central role of gender in social perception and cognition (16, 60), it may be unrealistic and impractical to fully "debias" or neutralize language in job postings: Instead, our approach advocates for moving toward neutrality. Similarly, it may not always be feasible to replace masculine language with perfectly synonymous gender-neutral terms, but again, our approach emphasizes striving for equivalence. Despite these limitations, the intervention is relatively light-touch and low-cost, making it an appealing first step for organizations beginning or seeking to amplify their gender diversity change efforts. It is important to emphasize, however, that this intervention is not a comprehensive solution to gender inequality, but rather, a starting point—a foot-in-the-door toward more substantive policies that dismantle biases within organizational systems (see (61, 62) for reviews), such as redesigning personnel processes to reduce bias (63–65). The idea that masculine defaults do not always have a gender-neutral equivalent is central to ultimately dismantling these norms (4). Organizations may need to undergo a broader value shift if they are truly committed to reducing gender inequality. More broadly, we encourage organizations and policymakers to view this intervention as a nudge to reflect on the language used in their job postings, especially when generating job postings is more systematized and automatic. This internal scrutiny could lead to a realignment between what organizations value, what they wish to convey, and how they communicate it more accurately, intentionally, and inclusively.

Conclusion

We demonstrate how a gender debiasing intervention that replaces masculine language with gender-neutral language creates more gender-diverse applicant pools by increasing application rates from women and men who do not "fit" with a strong male identity. Overall, these findings clarify when and how shifting masculine defaults can enhance diversity, inclusion, and belonging more broadly than anticipated.

Data, Materials, and Software Availability. Anonymized De-identified data have been deposited in osf (https://osf.io/mub7n/?view_only=65e0711f-34164331b38049ee5c2a5825) (36). Some study data available (De-identified data, materials, and code for analysis are available publicly. Raw data for Studies 1 and 2 containing identifying information of applicants (i.e., names) and job descriptions are not included for legal and proprietary reasons; the anonymized data are made public in our osf folder.)

1. NCSES, *Diversity and STEM: Women, Minorities, and Persons with Disabilities 2023* (National Science Foundation, Alexandria, VA, 2023). Special Report NSF 23-315. https://ncses.nsf.gov/wmpd.
2. U.S. Government Accountability Office, Financial services industry: Overview of representation of minorities and women and practices to promote diversity (GAO-23-106427). (U.S. Government Accountability Office, 2022). https://www.gao.gov/products/gao-23-106427.
3. E. Field *et al.*, Women in the workplace 2023. (*McKinsey & Company*, 2023), https://www.mckinsey.com/featured-insights/diversity-and-inclusion/women-in-the-workplace.
4. S. Cheryan, H. R. Markus, Masculine defaults: Identifying and mitigating hidden cultural biases. *Psychol. Rev.* **127**, 1022–1052 (2020).
5. F. Danbold, C. Bendersky, Balancing professional prototypes increases the valuation of women in male-dominated professions. *Organ. Sci.* **31**, 119–140 (2020).
6. F. Danbold, C. Bendersky, Perceived misalignment of professional prototypes reduces subordinates' endorsement of sexist supervisors. *J. Appl. Psychol.* **108**, 676–685 (2023).
7. A. Celani, P. Singh, Signaling theory and applicant attraction outcomes. *Pers. Rev.* **40**, 222–238 (2011).
8. B. L. Connelly, S. T. Certo, R. D. Ireland, C. R. Reutzel, Signaling theory: A review and assessment. *J. Manage.* **37**, 39–67 (2011).
9. D. Gaucher, J. Friesen, A. C. Kay, Evidence that gendered wording in job advertisements exists and sustains gender inequality. *J. Pers. Soc. Psychol.* **101**, 109–128 (2011).
10. J. B. Rosener, *America's Competitive Secret: Utilizing Women as a Management Strategy* (Oxford University Press, 1995).
11. J. C. He, S. K. Kang, Covering in cover letters: Gender and self-presentation in job applications. *Acad. Manage. J.* **64**, 1097–1126 (2021).
12. M. E. Heilman, Gender stereotypes and workplace bias. *Res. Organ. Behav.* **32**, 113–135 (2012), https://doi.org/10.1016/j.riob.2012.11.003.
13. L. A. Rudman, C. A. Moss-Racusin, J. E. Phelan, S. Nauts, Status incongruity and backlash effects: Defending the gender hierarchy motivates prejudice against female leaders. *J. Exp. Soc. Psychol.* **48**, 165–179 (2012).
14. R. J. Ely, D. E. Meyerson, An organizational approach to undoing gender: The unlikely case of offshore oil platforms. *Res. Organ. Behav.* **30**, 3–34 (2010).
15. S. Sczesny, M. Formanowicz, F. Moser, Can gender-fair language reduce gender stereotyping and discrimination? *Front. Psychol.* **7**, 25 (2016).
16. A. E. Martin, M. L. Slepian, The primacy of gender: Gendered cognition underlies the big two dimensions of social cognition. *Perspect. Psychol. Sci.* **16**, 1143–1158 (2020).
17. A. E. Martin, M. L. Slepian, Dehumanizing gender: The debiasing effects of gendering human-abstracted entities. *Pers. Soc. Psychol. Bull.* **44**, 1681–1696 (2018), https://doi.org/10.1177/0146167218774777.
18. M. E. Heilman, Sex bias in work settings: The lack of fit model. *Res. Organ. Behav.* **5**, 269–298 (1983).
19. B. Bhattacharyya, J. L. Berdahl, Do you see me? An inductive examination of differences between women of color's experiences of and responses to invisibility at work. *J. Appl. Psychol.* **108**, 1073–1095 (2023), https://doi.org/10.1037/apl0001072.
20. J. X. Goh *et al.*, Narrow prototypes and neglected victims: Understanding perceptions of sexual harassment. *J. Pers. Soc. Psychol.* **122**, 873–893 (2022).
21. J. K. Bosson, K. S. Michniewicz, Gender dichotomization at the level of ingroup identity: What it is, and why men use it more than women. *J. Pers. Soc. Psychol.* **105**, 425–442 (2013).
22. M. G. Nielson, K. M. Schroeder, C. L. Martin, R. E. Cook, Investigating the relation between gender typicality and pressure to conform to gender norms. *Sex Roles* **83**, 523–535 (2020).
23. T. E. Smith, C. Leaper, Self-perceived gender typicality and the peer context during adolescence. *J. Res. Adolesc.* **16**, 91–104 (2006).
24. J. A. Vandello, J. K. Bosson, D. Cohen, R. M. Burnaford, J. R. Weaver, Precarious manhood. *J. Pers. Soc. Psychol.* **95**, 1325–1339 (2008).
25. N. Ellemers, R. Spears, B. Doosje, Sticking together or falling apart: In-group identification as a psychological determinant of group commitment versus individual mobility. *J. Pers. Soc. Psychol.* **72**, 617–626 (1997).
26. R. R. Lau, Individual and contextual influences on group identification. *Soc. Psychol. Q.* **52**, 220 (1989).
27. T. Schmader, Gender identification moderates stereotype threat effects on women's math performance. *J. Exp. Soc. Psychol.* **38**, 194–201 (2002).
28. R. Spears, B. Doosje, N. Ellemers, Self-stereotyping in the face of threats to group status and distinctiveness: The role of group identification. *Pers. Soc. Psychol. Bull.* **23**, 538–553 (2016).
29. R. van Veelen, S. Otten, M. Cadinu, N. Hansen, An integrative model of social identification: Self-stereotyping and self-anchoring as two cognitive pathways. *Pers. Soc. Psychol. Rev.* **20**, 3–26 (2016).
30. J. L. Berdahl *et al.*, Work as a masculinity contest. *J. Soc. Issues* **74**, 422–448 (2018).
31. R. F. Baumeister, M. R. Leary, The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychol. Bull.* **117**, 497–529 (1995).
32. S. Cheryan, V. C. Plaut, P. G. Davies, C. M. Steele, Ambient belonging: How stereotypical cues impact gender participation in computer science. *J. Pers. Soc. Psychol.* **97**, 1045–1060 (2009).
33. A. B. Diekman, E. R. Brown, A. M. Johnston, E. K. Clark, Seeking congruity between goals and roles: A new look at why women opt out of science, technology, engineering, and mathematics careers. *Psychol. Sci.* **21**, 1051–1057 (2010).
34. C. Good, A. Rattan, C. S. Dweck, Why do women opt out? Sense of belonging and women's representation in mathematics *J. Pers. Soc. Psychol.* **102**, 700–717 (2012).
35. G. M. Walton, G. L. Cohen, A question of belonging: Race, social fit, and achievement. *J. Pers. Soc. Psychol.* **92**, 82–96 (2007).
36. J. C. He, S. Kang, De-biasing job ads by replacing masculine language increases gender diversity of applicant pools. OSF. https://osf.io/mub7n/?view_only=4c08e8c92c324b679f8ce8875cf02f0b. Deposited 1 June 2021.
37. C. Blevins, C. L. Mullen, Jane, John … Leslie? A historical method for algorithmic gender prediction *Digit. Humanit. Q.* **9**, (2015).
38. L. Mullen, gender: Predict gender from names using historical data (R package version 0.5.3, 2019). https://github.com/lmullen/gender. Accessed 28 January 2021.
39. J. Pilcher, Names and "doing gender": How forenames and surnames contribute to gender identities, difference, and inequalities. *Sex Roles* **77**, 812–822 (2017).
40. P. G. Erwin, A review of the effects of personal name stereotypes. *Represent. Res. Soc. Psychol.* **20**, 41–52 (1995).
41. D. N. Figlio, Boys named sue: Disruptive children and their peers. *Educ. Finance Policy* **2**, 376–394 (2007).
42. H. Barry Jr., A. S. Harper, Increased choice of female phonetic attributes in first names. *Sex Roles* **32**, 809–819 (1995).
43. J. J. Endendijk, M. G. Groeneveld, J. Mesman, The gendered family process model: An integrative framework of gender in the family. *Arch. Sex. Behav.* **47**, 877–904 (2018).
44. G. M. Alexander, K. John, T. Hammond, J. Lahey, Living up to a name: Gender role behavior varies with forename gender typicality. *Front. Psychol.* **11**, 604848 (2021).
45. J. W. Pennebaker, R. J. Booth, M. E. Francis, Linguistic inquiry and word count: LIWC. (LIWC, 2007).
46. J. A. Bartz, J. E. Lydon, Close relationships and the working self-concept: Implicit and explicit effects of priming attachment on agency and communion. *Pers. Soc. Psychol. Bull.* **30**, 1389–1401 (2004).
47. S. L. Bem, The measurement of psychological androgyny. *J. Consult. Clin. Psychol.* **42**, 1–5 (1974).
48. C. Hoffman, N. Hurst, Gender stereotypes: Perception or rationalization? *J. Pers. Soc. Psychol.* **58**, 197–208 (1990).
49. L. A. Rudman, S. E. Kilianski, Implicit and explicit attitudes toward female authority. *Pers. Soc. Psychol. Bull.* **26**, 1315–1328 (2000).
50. V. K. Gupta, D. B. Turban, N. M. Bhawe, The effect of gender stereotype activation on entrepreneurial intentions. *J. Appl. Psychol.* **93**, 1053–1061 (2008).
51. A. T. Jebb, L. Tay, Introduction to time series analysis for organizational research. *Organ. Res. Methods* **20**, 61–94 (2017).
52. A. T. Jebb, L. Tay, W. Wang, Q. Huang, Time series analysis for psychological research: Examining and forecasting change. *Front. Psychol.* **6**, 767 (2015), https://doi.org/10.3389/fpsyg.2015.00727.
53. W. Newey, K. West, A simple, positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix. *Appl. Econom.* **33**, 125–132 (2014).
54. D. W. K. Andrews, J. C. Monahan, An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica* **60**, 953 (1992).
55. A. Zeileis, Econometric computing with HC and HAC covariance matrix estimators. *J. Stat. Softw.* **11**, 1–17 (2004).
56. K. Deaux, B. Major, Putting gender into context: An interactive model of gender-related behavior. *Psychol. Rev.* **94**, 369–389 (1987), https://doi.org/10.1037/0033-295X.94.3.369.
57. E. J. Castilla, H. J. Rho, The gendering of job postings in the online recruitment process. *Manage. Sci.* **69**, 6417–7150 (2023), https://doi.org/10.1287/mnsc.2023.4674.
58. P. Parasurama, J. Sedoc, A. Ghose, Gendered information in resumes and hiring bias: A predictive modeling approach. SSRN [Preprint] (2022). https://doi.org/10.2139/ssrn.4074976 (Accessed 24 May 2024).
59. A. H. Eagly *et al.*, Gender stereotypes have changed: A cross-temporal meta-analysis of U.S. public opinion polls from 1946 to 2018. *Am. Psychol.* **75**, 301–315 (2020).

60. A. B. Diekman, T. Schmader, "Gender as embedded social cognition" in *The Oxford Handbook of Social Cognition*, D. E. Carlston, K. Hugenberg, K. L. Johnson, Eds. (Oxford University Press, ed. 2, 2024), pp. 471–501.
61. I. N. Onyeador, S. Mobasseri, H. L. McKinney, A. E. Martin, A future for organizational diversity training: Mobilizing diversity science to improve effectiveness. *Acad. Manage. Perspect.* **38**, 392–414 (2024).
62. N. M. Stephens, L. A. Rivera, S. S. M. Townsend, The cycle of workplace bias and how to interrupt it. *Res. Organ. Behav.* **40**, 100137 (2020).
63. I. Bohnet, A. van Geen, M. Bazerman, When performance trumps gender bias: Joint vs. separate evaluation. *Manage. Sci.* **62**, 1225–1234 (2016).
64. E. H. Chang, E. L. Kirgios, A. Rai, K. L. Milkman, The isolated choice effect and its implications for gender diversity in organizations. *Manage. Sci.* **66**, 2291–2799 (2020).
65. J. C. He, S. K. Kang, N. Lacetera, Opt-out choice framing attenuates gender differences in the decision to compete in the laboratory and in the field. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2108337118 (2021), https://doi.org/10.1073/pnas.2108337118.