Beliefs about minority representation in policing and support for diversification

Kyle Peyton[a,1], **Chagai M. Weiss**[b,c], **and Paige Vaughn**[d]

[a]Institute for Humanities and Social Sciences, Australian Catholic University; [b]Department of Political Science, University of Wisconsin – Madison; [c]The Middle East Initiative at Harvard Kennedy School; [d]Department of Sociology and Criminology, Spring Hill College

This manuscript was compiled on July 7, 2022

Diversification has been widely promoted as a reform for reducing racial disparities in police-civilian interactions and increasing police legitimacy. Recent evidence suggests civilians are more likely to trust police forces that demographically represent the populations they serve, and that diversity enhancing reforms can improve police treatment of minority communities. However, nearly every municipal police department in the United States is predominately White and male, and the scale of minority under-representation in most departments is substantial. The widespread implementation of diversity reforms likely requires awareness of existing disparities and support for policy change; yet little is known about beliefs and preferences among police or the general public. Here we investigate beliefs about minority representation and preferences for diversification using a series of experiments, including on civilians and police from a city with one of the largest demographic disparities in the country. We find that the general public overestimates minority representation in policing, and that correcting this unfounded optimism with accurate information decreases trust in police. These information interventions also increase support for hiring decisions that favor minority applicants, and willingness to vote for local policy change. Additional paired decision-making experiments demonstrate that, even without these corrections, both current officers and the residents they police prefer hiring new officers from under-represented groups, independent of civil service exam performance and other criteria. Overall, these findings suggest that neither the attitudes and preferences of police or the general public pose a major barrier to diversification.

policing | diversity | representation | bureaucracy

Repeated instances of police violence against unarmed civilians in the United States have drawn widespread attention to long-standing concerns about racially biased policing, and renewed interest in various reforms aimed at improving police-community relations (1–4). In addition to community policing (2, 5), body-worn cameras (6, 7), and officer training initiatives (8, 9), police department diversification has been widely promoted as a policy tool for improving police-community relations and promoting just and equitable policing (3, 10). A growing body of research indicates that diversification is associated with numerous benefits, including greater trust and cooperation (11, 12), increased crime reporting (13), reduced use of force against civilians, and improved treatment of minority communities (3).

Despite these potential benefits, and long-standing calls from policy-makers to make police forces demographically representative of their communities (15, 16), non-White officers are severely under-represented in most police departments (see Figure 1). Pooling across 474 large departments – those employing at least 100 officers – we find that approximately 62% of officers are White, compared with 44% of civilians in the communities they police (14, 16). The difference between the share of non-White residents and non-White officers exceeds 20 percentage points in 60% of these departments. Recent estimates for the largest 97 departments – representing more

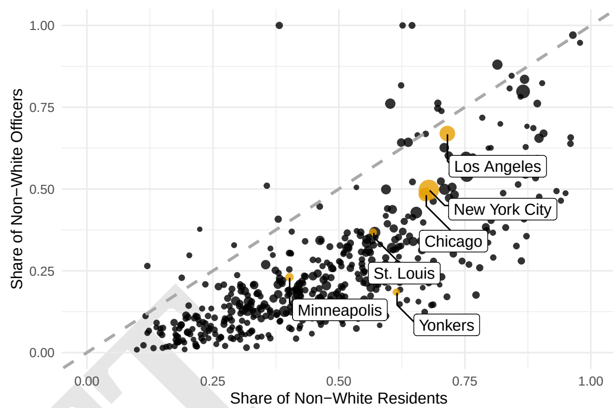

Fig. 1. The share of non-White officers (vertical axis) compared with the share of non-White residents (horizontal axis) in 474 local agencies that employ at least 100 officers. Each point on the graph represents a jurisdiction/department, with size proportional to the size of the resident population. Points below the gray line denote police departments that under-represent the communities they serve (approx. 95% of departments). Officer demographics come from the most recent Law Enforcement Management and Administrative Statistics (LEMAS) survey (2016), which sampled all local agencies that employed at least 100 officers with certainty (14). Estimates of the demographic proportions for the resident population in each jurisdiction come from the U.S. Census. Together, these data cover 242,240 officers from jurisdictions with a total civilian population of more than 112 million.

Significance Statement

Diversification is often proposed as a policy to promote more equitable policing. Yet little is known about preferences for police diversification among officers or the general public. We show that the general public greatly overestimates minority representation in policing, and that information interventions correcting these misperceptions increase support for diversification, heighten respondents' willingness to vote for policy change, and reduce trust in the police. Additional experiments demonstrate that both civilians and officers prefer increased minority representation in the absence of any corrective information. These findings suggest that contemporary attitudes among voters and officers are not a major barrier to police diversification.

K.P., C.W., and P.V. designed research; K.P. performed research and project administration; K.P. and C.W. analyzed data; K.P., C.W., and P.V. wrote the paper.

The authors declare no conflict of interest.

[1]To whom correspondence should be addressed. Email: kyle.peyton@acu.edu.au.

than a third of all local police in the U.S. – reach similar conclusions: 56% of officers are White, compared to 36% of the civilians in their jurisdictions (17). Analyses of historical data also suggest that most departments have become even less representative of the communities they serve over the past decade (16).

The scale and persistence of minority under-representation in U.S. policing suggests the need for reforms that explicitly target the hiring and recruitment process. There are, however, at least two potential political challenges. First, given the enduring political salience of crime and policing, public opinion – or policy makers' beliefs about public opinion – tends to shape the direction of potential reforms as well as the scope of policy change (4, 18). Second, police officers – through unions and professional organizations – exert power over policy outcomes at both the local and national level. Even reforms that enjoy broad public support, such as community policing and body worn cameras, can face implementation challenges without adequate "buy-in" from police officers (19, 20). In short, successfully implementing diversity reforms across numerous U.S. municipalities will require public awareness of existing disparities, and support for policy change. Yet little is known about attitudes toward police diversification among the general public, or among police officers.

Here, we investigate beliefs about minority representation in policing and attitudes toward diversification using a series of surveys and experiments fielded across three different samples: a national sample of U.S. adults, a municipal sample of Yonkers, NY residents, and a police sample of sworn officers from the Yonkers Police Department (YPD). These paired samples of police and residents provide a unique opportunity to study attitudes towards diversification in a jurisdiction with one of the least representative police forces in the country (see Figure 1). We use these data to shed new light on three important questions. First, is the general public aware of the lack of diversity in U.S. policing? Second, does the provision of accurate information about minority under-representation affect public support for police diversification? Third, are the hiring preferences of current police officers, and community residents, affected by the race/ethnicity and gender of potential police recruits?

Beliefs about minority representation in policing

Prior research demonstrates that beliefs about progress toward equity and inclusion in the United States are overly optimistic, especially in the domain of racial economic inequality (21–23). For example, recent data show that a majority (> 60%) of U.S. adults underestimate Black-White wealth inequality by at least 20 percentage points (22). Similar patterns also hold for beliefs about residential segregation and economic mobility (24, 25).

Given this prior work, and the fact that policing data are both notoriously scattered and infrequently publicized (17, 26), we anticipated that most individuals would have inaccurate beliefs about minority representation in U.S. policing. While official statistics on the race and gender composition of police departments are available from the Law Enforcement Management and Administrative Statistics (LEMAS) survey, we are unaware of any prior work on public *perceptions* of police force demographics. It was therefore unclear whether beliefs about police force diversity would be overly optimistic, or too pessimistic.

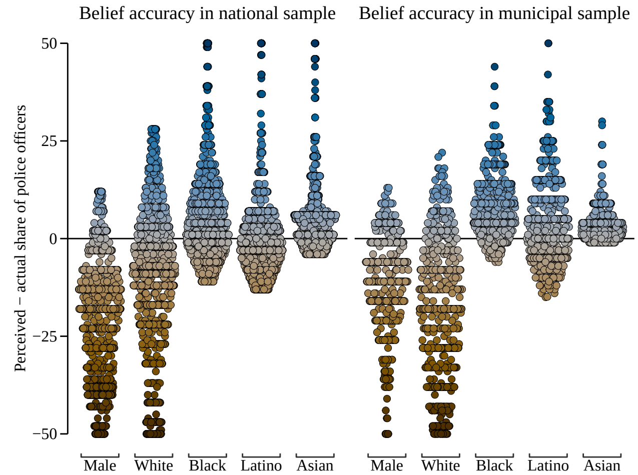

Fig. 2. Differences between perceived and actual shares of police officers in the national (left) and municipal (right) samples. Each point on the graph represents the difference between an individual's best guess about the percentage of police officers in the United States (Yonkers, NY) that belong to each group and the actual percentages according to official statistics. Points are jittered to avoid over-plotting and shaded so that darker blue (brown) colors denote greater levels of over-(under-) estimation.

We elicited public perceptions of police force diversity in a national survey of U.S. adults fielded in July 2021 (N = 2,017), and in the second wave of a municipal panel survey of Yonkers, NY residents fielded in October 2021 (N = 644). Respondents in the national (municipal) samples were asked to provide their best guess of the share of police officers in the United States (Yonkers, NY) from each of the four race/ethnicity (Black, White, Hispanic/Latino, and Asian) and two gender (male or female) groups for which official statistics are available. Given that individuals tend to overestimate the size of minority groups (27, 28), we followed prior work (4, 29) and provided respondents with the relevant population shares for each (e.g., "19% of Yonkers residents are Black"). Each was presented in randomized order, and responses were required to add to 100% across the four race/ethnicity measures, as well as the binary gender measure. We provide detailed descriptions of survey design, recruitment procedures, sample characteristics, and question wordings in Supporting Information (SI) Appendix S1.

Results. Figure 2 shows the differences between each respondent's estimate for a given group and the actual share among police officers in the United States (left panel) and Yonkers (right panel). Positive (negative) values denote over- (under-) estimation of the true share. This provides clear descriptive evidence that beliefs about minority representation in policing are overly optimistic, regardless of whether individuals were making inferences about U.S. police in aggregate (national sample), or their local police department (municipal sample).

On average, respondents underestimated the share of male officers by 22 percentage points ($\widehat{se} = 0.37$) in the national sample and 19 percentage points ($\widehat{se} = 0.32$) in the municipal sample. Likewise, beliefs about the percentage of White officers were roughly 14 percentage points ($\widehat{se} = 0.32$) too low at the national level, and roughly 19 percentage points ($\widehat{se} = 0.70$) too low at the municipal level. In both samples, the majority of respondents underestimated the White (male) officer shares by more than 10 percentage points. Conversely, respondents overestimated minority representation: Black, Hispanic/Latino, and Asian officers, as well as female officers, were perceived to be significantly more prevalent than they actually were in both samples.

In SI Appendix S2.1.8 we investigate whether certain groups of respondents (e.g., Whites, Republicans) are more likely to hold incorrect beliefs, or have more extreme beliefs. We find some evidence that misperceptions are correlated with respondents' background characteristics, but these associations are weak and inconsistent across measures. We find stronger evidence that beliefs about minority representation are correlated across domains; for example, respondents' misperceptions about gender diversity are a better predictor of their misperceptions about racial diversity than their background characteristics.*

Effects of information interventions on attitudes and behavior

Biased beliefs have important implications for politics and policy: aggregate preferences (and policy outcomes) in an uninformed electorate can be radically different from one in which individuals are adequately informed (30, 31). A key implication for the present research is that overly optimistic beliefs about police diversity may constrain public support for policy change, which could partly explain the scale and persistence of minority under-representation in policing. Here, we examine the causal link between belief accuracy and support for diversification using randomized experiments that provide accurate information about minority representation in policing.

An important advantage of information provision experiments in general is that they can be used to test for causal links between belief accuracy and other outcomes without deception (32). A growing body of empirical research also supports their efficacy: across a variety of contexts, individuals typically update their beliefs in the direction of the evidence they receive (33–37). The belief changes induced by information provision experiments do not, however, always have downstream effects on relevant attitudes and behaviors (35, 36). Given the absence of similar work on minority representation in policing, it was unclear whether information interventions would change beliefs, or have any downstream effects on support for diversification.

Related research on representative bureaucracy suggests that, under some circumstance, minority representation can influence public trust and willingness to cooperate with police and other "street-level bureaucrats". But this work, which draws primarily on cross-sectional surveys and vignette experiments about hypothetical agencies, has reached mixed conclusions (12, 38–40). Moreover, changing individuals' trust in government does not necessarily lead to downstream effects on policy preferences (41), and members of majority groups (e.g., White voters) are often opposed to policies that seek to increase minority representation (42–44).

To measure the effects of providing information about police officer diversity on attitudes and behaviors, we embedded information provision experiments in our national survey of the U.S. adult population and our follow-up survey of Yonkers residents (see SI Appendix S1.3 for design details; SI Appendix S3 for pre-registration). After eliciting respondents' beliefs about police officer diversity (see Fig. 2), they were randomly assigned to receive accurate information about police officer diversity at the national (municipal) level alongside the estimates they previously provided (treatment group). Those that were instead assigned to a no information condition did not receive this information (control group).

Respondents in the national sample (N = 2,017) were assigned to two additional conditions, one that included a description of findings from a recent study demonstrating the positive effects of police diversification (3); and another that provided this description alongside the accurate information about officer diversity. We did not detect any differences between the treatment effects for the information only condition and the treatment effects for either of these additional conditions (see SI Appendix S2.1.5). In anticipation of the smaller sample size in the municipal sample (N = 644), we did not include these additional treatment arms.

Here, we focus on the effects that correcting misperceptions about minority representation in policing – via the provision of accurate information – have on four attitudinal outcomes and two behavioral outcomes. Our primary attitudinal outcomes of interest (measured in both experiments) capture stated support for implementing affirmative action programs to increase recruitment and hiring of police officers from minority groups, and preferences for tie-breaking hires in favor of minority applicants. In the municipal sample, we also included measures of trust and confidence in the police (2-item index, $\alpha = 0.83$), and willingness to cooperate with police (4-item index, $\alpha = 0.73$). These indices were constructed using items that regularly appear in surveys of civilian attitudes toward police (1, 2, 45, 46).

Support for affirmative action was measured using a 4-item index of support for programs targeting each minority group: "Female officers", "Black officers", "Hispanic/Latino officers", and "Asian officers" (each presented in random order, using a 7-point scale with a neutral midpoint; $\alpha = 0.98$ in municipal sample, $\alpha = 0.96$ in national sample). Support for tie-breaking hires was also measured using a 4-item index of respondents' preferred option for deciding between "two equally qualified applicants for police officer" (each decision presented in random order; $\alpha = 0.89$ in municipal sample, $\alpha = 0.81$ in national sample). For each comparison, respondents chose between hiring the minority applicant (e.g. the "Black applicant"), coded 1; the non-minority applicant (e.g., the "White applicant") coded -1 or a third option of "Random selection (e.g., let a coin flip decide)", coded 0.

Finally, we included two behavioral outcomes in the municipal survey. The first, inspired by recent information experiments on racial discrimination (35), provided individuals with an opportunity to donate real money to a local non-profit that works to support Black individuals in law enforcement. For this outcome, all respondents were entered into a \$50 raffle (with a 1 in 20 chance of winning) and decided whether to keep this money versus make a real donation to the Black officers organization. We also provided individuals with an opportunity to cast a vote in favor of one of four police reforms: civilian oversight, diversification, community policing,

*For example, the R^2 from a linear regression of respondents' belief accuracy for the White officer share on their partisanship, race/ethnicity, education, and sex is less than 0.04 in both samples. By comparison, the R^2 from a linear regression of respondents' belief accuracy for the White officer share on belief accuracy for the male officer share is greater than 0.10 in both samples.

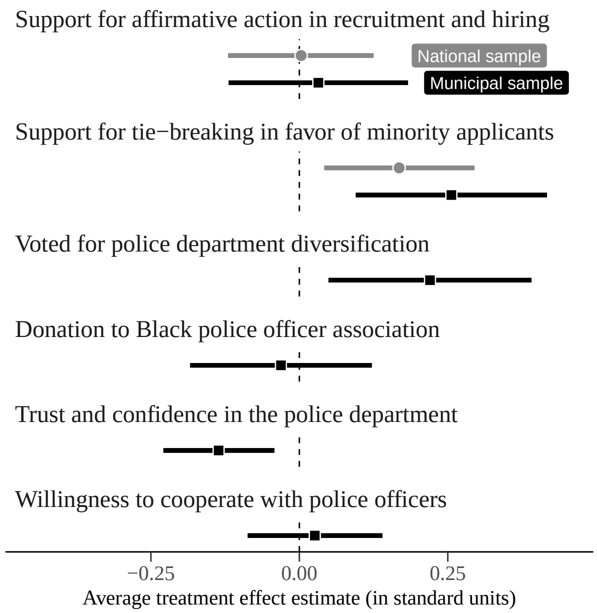

Fig. 3. Estimated treatment effects of accurate information about minority representation in policing on attitudes and behaviors in the national (grey) and municipal (black) samples. Treatment effects were estimated using linear regression of the outcome on treatment assignment, with standard errors (and 95% confidence intervals) based on HC2 robust standard errors. To facilitate comparisons, all estimates are standardized using using Glass's Δ, which scales outcomes by the standard deviation in the control group.

or body worn cameras. Each of these reforms were being actively discussed between YPD leadership and Yonkers residents at various community meetings that took place while the municipal surveys were in the field. A detailed description of each reform was provided to respondents during the survey, and they were instructed that the votes would be tallied and presented to the mayor and police commissioner in aggregate anonymized form. SI Appendix S1.3.1 provides additional details about outcome measurement, including question wordings and response categories for each survey item.

Results. Figure 3, shows the average causal effect of the information interventions on each of the six outcome measures previously described. Effects were estimated using linear regression of the outcome on treatment assignment, with standard errors (and 95% confidence intervals) based on HC2 robust standard errors. To facilitate comparisons, all estimates are standardized using using Glass's Δ, which scales outcomes by the standard deviation in the control group.

First, we find that the effect on support for affirmative action programs was statistically indistinguishable from zero in both the national ($\delta = 0.00$, $\widehat{se} = 0.06$, $t = 0.05$, $P = 0.96$) and municipal samples ($\delta = 0.03$, $\widehat{se} = 0.08$, $t = 0.42$, $P = 0.68$). However, we do find significant positive effects on preferences for tie-breaking hires in favor of minority group applicants competing with "equally qualified" majority group applicants (national sample: $\delta = 0.17$, $\widehat{se} = 0.06$, $t = 2.61$, $P < 0.01$; municipal sample: $\delta = 0.26$, $\widehat{se} = 0.08$, $t = 3.13$, $P < 0.01$). For context, these effect sizes are larger than the average differences between untreated White and non-White respondents (0.09 in national sample; 0.13 in municipal sample).

These results suggests that generic support for affirmative action may be more resistant to change than preferences for specific policy implementations (e.g., tie-breaking in favor of under-represented groups). One possible explanation for this apparent disconnect is that public support for a given policy is often shaped by perceptions of that policy's substantive implications (4, 47). Prior work finds that Americans do not have a coherent understanding of what "affirmation action" actually means, and that beliefs in prevailing myths (e.g., that it is a quota system) are strongly correlated with opposition (48). Related work in political psychology has also found that those who oppose affirmative action in the abstract do not necessarily oppose specific applications, including tie-breaking (49, 50).

SI Appendix S2.1.1. provides supplementary analyses that estimate effects on each index component (e.g., support for programs targeting "Black officers"). These estimates are not statistically distinguishable from one another, suggesting the precision gains from dimension reduction are worth the potential drawback of using summary indices that abstract away from variation in attitudes towards each group.

Turning to the additional outcomes measured in the municipal sample, we find that the effect on donations to the Black officers association was not statistically distinguishable from zero ($\delta = -0.03$, $\widehat{se} = 0.08$, $t = 0.39$, $P = 0.69$). Here, roughly 56% of treated respondents agreed to donate some amount (avg. donation: \$17.70), versus 57% in the control group (avg. donation: \$18.40). A recent study on racial discrimination in the labor market found similar results: information interventions improved belief accuracy but did not increase donations to a civil rights group (35).

The intervention here did, however, cause a significant increase in residents' willingness to vote in favor of diversifying their local police department ($\delta = 0.22$, $\widehat{se} = 0.09$, $t = 2.53$, $P = 0.01$). For context, this effect size translates to a difference of about 8 percentage points on a binary scale: 13% of respondents in the control group voted for diversification compared to 21% in the treatment group. One possible explanation for these effects is that individuals view police diversification as a local policy issue that should be addressed by municipal government (i.e., the mayor and police commissioner), rather than a non-profit.

Finally, we find that providing information about the (lack of) diversity at YPD caused a significant decrease in trust in the police ($\delta = -0.14, \widehat{se} = 0.05$, $t = 2.84$, $P < 0.01$). For context, this effect size was approximately 8 percentage points when measured on the same (single item) scale used in Gallup's national survey of confidence in institutions, which found a 5 percentage point decrease in trust following the murder of George Floyd in May 2020 (46). Despite this significant negative effect on Yonkers residents' trust in YPD, we find the effect on their willingness to cooperate with police officers was not statistically distinguishable from zero ($\delta = 0.03, \widehat{se} = 0.06$, $t = 0.45$, $P < 0.01$).

SI Appendix S2.1.2 compares the estimates reported here with covariate-adjusted estimates. We find limited precision gains from regression-adjustment in this application. We also report supplementary analyses for effect heterogeneity as a function of pre-treatment covariates (including partisanship, race/ethnicity, and belief accuracy) in SI Appendix S2.1.3. These analyses do not reliably identify sub-groups for which stronger (or weaker) causal effects are obvious.

SI Appendix S2.1.4-S2.1.6 provides supplementary analyses to explore potential alternative mechanisms which might explain the results from the information provision experiments. Overall, we find compelling evidence that information provision increased support for diversification (and reduced trust) via belief updating, rather than by causing individuals to attach more importance to the issue of minority representation in policing. For example, we find that exposure to high-quality research on the benefits of police diversification did not lead to attitude change unless also paired with the information interventions described here.

Effects of race and gender on hiring preferences of local residents and police

Direct discrimination – explicitly different treatment of individuals caused by their group membership – has been identified as an important disparity generating mechanism across a wide range of contexts in the United States (51). The results from the previous section demonstrate that correcting unfounded optimism about minority representation in policing can increase public support for tie-breaking hires in favor of minority applicants, as well as local residents' willingness to vote for police department diversification. Our interpretation is that factual information affected these outcomes by reducing gaps between perceptions and reality.

This suggests public support for diversification is not necessarily constrained by underlying preferences for White (male) over non-White (female) officers. However, the information experiments do not directly identify how a minority applicant's race/ethnicity (or gender) might affect the likelihood they would be hired by a police department. If, for example, community residents or other police officers within a department have preferences that systematically favor hiring White (male) officers then minority under-representation in policing could be explained, at least in part, by direct discrimination.

Conjoint experiments have been widely used to study the role that direct discrimination plays in contexts involving multidimensional choices, and they offer several advantages (52–56). First, the randomization of multiple attributes allows us to estimate the marginal effects of applicant race/ethnicity and gender alongside other factors that are heavily weighted in police recruitment policies, such as civil exam scores and residency requirements. The design also better reflects the multidimensional nature of the decision making task, and prior research has found strong correspondence between hypothetical choices in conjoint experiments and real world behavior (57). Finally, conjoint designs are less susceptible to social desirability biases than vignette designs because they randomize sensitive features (e.g., race) alongside other relevant attributes (54, 58).[†]

To measure how the hiring preferences of police officers and civilians are affected by the race/ethnicity and gender of potential police recruits, we embedded a police recruitment conjoint experiment in the first wave of a municipal panel survey of Yonkers residents in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130) and a survey of Yonkers police officers in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500). We provide a detailed description of this experimental design in SI Appendix S1.4 (see S1.2 for recruitment procedures and sample characteristics; S4 for pre-registration).

A unique advantage in the present context is that we can examine whether the preferences of YPD officers differ systematically from Yonkers residents. To our knowledge, this is the first attempt at directly estimating how race/ethnicity and gender of applicants affect the hiring preferences of police officers and community residents. In both our samples, respondents made choices between potential recruits to the YPD that varied independently across their age, race/ethnicity, sex, civil service exam performance, education, prior occupation, length of municipal residency, and their motivation for applying to become a police officer. Attribute levels were chosen based on a combination of interviews with YPD officers, historical data on officer applicants, and prior survey work on police officers' motivations and background characteristics (60, 61).[‡]

In practice, exam scores and residency requirements receive disproportionate weight in the hiring process. This is because municipal police departments require all applicants to complete a civil service exam, and those who pass are then rank-ordered by their test score on an "eligibility list". This is typically the first formal stage of the hiring process, and only those on the list are eligible to proceed to subsequent stages (physical fitness tests, background investigations, oral interviews, etc.). Many departments also impose a residency requirement; for example, that potential applicants must live within a certain distance of the city for a period of at least 3 months immediately preceding the exam (YPD's policy). Therefore, residency and exam performance are among the primary hiring criteria at municipal police departments. For example, if there are 100 applicants on the eligibility list and 30 openings then, all else equal, the 30 with the highest exam scores will be selected.

Results. We estimate the Average Marginal Component Effects (AMCEs) of randomly assigned attributes on the probability of selecting an applicant (binary outcome) using linear regression, with robust standard errors clustered at the respondent level to correct for within-respondent clustering. Here, we focus on the marginal effects that an applicant's race/ethnicity, sex, civil service exam performance, and length of residency have on their probability of being selected (see SI Appendix S2.2 for estimated AMCEs and marginal means of all randomized attributes). Figure 4 shows the estimated effects that each of these characteristics have (relative to the reference category) on the probability of selecting an applicant for hiring in the resident (black) and police officer (blue) samples.

As expected, higher performance on the civil service exam and longer residency have large causal effects on the probability that a given applicant is selected. For example, the effect of scoring in the Top 1% on the civil service exam (relative to the Top 25%) is 0.15 ($\widehat{se} = 0.01$, $t = 11.28$, $P < 0.01$) in the civilian sample and 0.28 in the police sample ($\widehat{se} = 0.03$, $t = 8.94$, $P < 0.01$). The effect of being a long-term community

[†] We also followed best practices to mitigate these potential threats by conducting anonymous online surveys, and providing additional assurances of anonymity and data privacy to participants at the beginning of the survey. In general, the potential threat that survey respondents systematically engage in "preference falsification" and do not provide honest answers to sensitive questions seems low. For example, a recent meta-analysis covering 30 years of research found, if anything, survey respondents slightly over-report racist attitudes (59).

[‡] To avoid implausible cases (e.g., school teacher's with GED's) we employed restricted randomization on the education and occupation attributes such that potential applicants that were previously school teachers or social workers always had education levels of at least a Bachelor's degree or higher. All estimates are adjusted to account for this conditional independence, which is a common feature in conjoint experiments that seek to avoid generating implausible profiles (62).

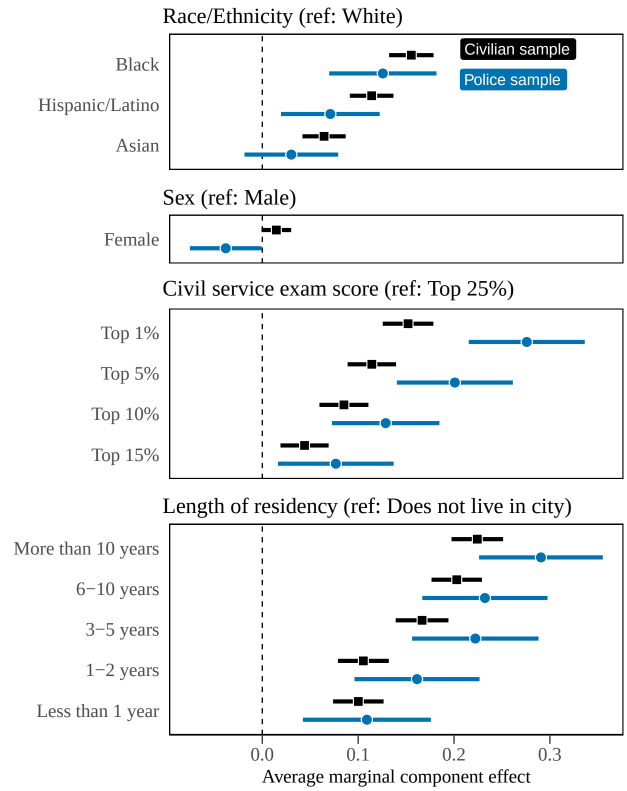

Fig. 4. Average Marginal Component Effects (AMCEs) of randomly assigned characteristics of police officer applicants on probability of selecting an applicant for hiring. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Estimates are adjusted to account for randomization constraints on the education and occupation attributes. Civilian sample (black): municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample (blue): survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

resident (i.e., more than 10 years) is 0.22 ($\widehat{se} = 0.01$, $t = 16.34$, $P < 0.01$) in the civilian sample and 0.29 ($\widehat{se} = 0.03$, $t = 8.84$, $P < 0.01$) in the police sample. The large between-sample differences at the top of the score distribution suggests that officers afford more weights to high scoring applicants than community residents. None of the estimated AMCEs for length of residency were statistically distinguishable between samples.

Independent of these other relevant characteristics, both YPD officers and community residents clearly prefer non-White over White police recruits. On average, White applicants were selected for hiring with probability 0.42 in the civilian sample and 0.45 in the police sample.[§] An application from a Black (relative to White) individual causes an increase in the probability of selection by 0.16 (i.e., 16 percentage points) in the civilian sample ($\widehat{se} = 0.01$, $t = 13.14$, $P < 0.01$) and 0.13 in the police sample ($\widehat{se} = 0.03$, $t = 4.40$, $P < 0.01$). Despite large demographic differences between samples (police: 82% White, 85% male; civilian: 45% White, 41% male) none of the estimated AMCEs for race/ethnicity were statistically distinguishable.[¶]

When considering prospective applicants' gender, however, female applicants do not appear to have a systematic advantage over males. The estimated AMCE for female, relative to male, applicants corresponds to an increase in the probability of hiring of just 0.01 (i.e., about 1 percentage point) in the civilian sample ($\widehat{se} = 0.01$, $t = 1.86$, $P = 0.06$), and a small *decrease* of 0.04 (i.e., about 4 percentage points) in the police sample ($\widehat{se} = 0.02$, $t = 2.00$, $P = 0.05$). This suggests that female applicants, on average, may be slightly disadvantaged relative to male applicants.

We explore causal interactions among race/ethnicity, sex, and exam performance in SI Appendix S2.2.2. We find that bias against female (v. male) applicants appears unique to White females in the police sample, whereas there is no evidence of bias against non-White females in either sample (Fig. S40). We also estimate causal interactions between exam scores and race/ethnicity (Fig. S41-S45), as well as exam scores and sex (Fig. S46-47). These results suggest non-White applicants are preferred at every level of exam performance. Moreover, non-White applicants with lower test scores are, all else equal, preferred to White applicants with higher scores. We find minimal differences between samples.

To illustrate the substantive implications of these results, we can estimate predicted probabilities for different types of applicants that only vary on race/ethnicity and gender. For example, consider a 27 year old Black male applicant with a high school education, who has lived in Yonkers for 10+ years, previously worked as a security guard, scored within the Top 25% on the exam, and listed "helping people" as their primary motivation. This applicant would be selected with probability 0.70 ($\widehat{se} = 0.02$) in the civilian sample and 0.70 in the police sample ($\widehat{se} = 0.06$).

On the other hand, an otherwise similar White male applicant would be selected with probability 0.56 ($\widehat{se} = 0.07$) among police and 0.54 ($\widehat{se} = 0.03$) among residents. White female applicants fare similarly, with selection probabilities of 0.52 ($\widehat{se} = 0.07$) among police and 0.56 ($\widehat{se} = 0.03$) among residents. A Black female applicant with the same characteristics would be selected with probability 0.66 ($\widehat{se} = 0.07$) by police and 0.72 ($\widehat{se} = 0.02$) by residents. Overall, the results from these paired decision making experiments suggest a remarkable degree of similarity between police and public preferences for minority hires.

Discussion

Despite long-standing normative concerns about minority under-representation in policing (15, 16), and a growing body of empirical evidence documenting the potential benefits of diversification (3, 11–13), most U.S. police departments are dominated by White men. The scale and persistence of minority under-representation suggests the need for reforms that increase hiring and recruitment from under-represented groups; yet little is known about support for diversification among police or the general public. The results reported here shed new light on support for police diversification across multiple samples, including the general public and sworn officers from a large urban police department.

Consistent with recent work on beliefs about the scale and persistence of Black-White inequality in the U.S. (21–23), we

[§] Given that respondents must always choose between two potential recipients, the expected value is 0.50 under the null hypothesis of indifference.

[¶] SI Appendix S2.2.3 explores heterogeneity by covariates (e.g., race/ethnicity and partisanship) among Yonkers residents. Although Non-White respondents (as well as females and Democrats) were significantly more supportive of minority applicants, we did not identify any sub-groups that disfavored minority applicants.

find clear evidence of unfounded optimism about minority representation in policing. This holds regardless of whether individuals are making inferences about U.S. police in aggregate, or their local police department. We also find that correcting these biased beliefs has downstream consequences. While reducing the gap between perceptions and reality decreased trust in the police, it also caused an increase in support for hiring decisions that favor minority applicants, as well as local residents' willingness to vote for diversification over other police reforms. Extending fundamental insights about the political implications of biased beliefs (30, 31), this suggests unfounded optimism about police diversity may constrain public support for policy change.

These results are particularly noteworthy given that belief updating does not necessarily lead to attitude change (33–37), and opposition to policies aimed at increasing minority representation is the norm in other contexts (42–44). The observation that preferences for specific hiring policies favoring minority applicants were less resistant to change than generic support for affirmative action also underscores the utility of direct questioning (4, 47, 49, 50). More broadly, the finding that accurate information about the (lack of) minority representation in policing decreased public trust – but also increased support for policy change – suggests limits to normative perspectives that emphasize the value of trust in the police as an end in itself.

Overall, these information experiments demonstrate that exposure to accurate information about minority underrepresentation can increase public support for diversification by reducing unfounded optimism about officer diversity. Our interpretation is that demand for increased minority representation already exists, and information exposure increases support by correcting biased beliefs. Consistent with this interpretation, our paired experiments in Yonkers, NY demonstrate that – even in the absence of any corrective information – both current officers and community residents prefer hiring new officers from under-represented groups, independent of civil service exam performance and other criteria.

These paired experiments provide unique insights about preferences among both officers and residents in a jurisdiction with one of the least representative police forces in the country. For example, 78% of YPD officers are White compared to 34% of the adult population: a difference ($\sim$ 44 percentage points) that is more extreme than 92% of the jurisdictions where official statistics are available. Police-community relations in Yonkers have also suffered from a long history of conflict and distrust; including, for example, a 2007 investigation by the U.S. Department of Justice into allegations of discriminatory policing that took nearly a decade to resolve.[‖]

Although it would be premature to conclude that officers and residents across thousands of other local law enforcement jurisdictions have similar preferences, there are few places where representational disparities might suggest a sharper divide between police and the public. Taken together, these findings suggest that neither the attitudes or preferences of officers or the general public pose a major barrier to police diversification.

[‖] see: https://www.justice.gov/opa/file/910316/download, and https://www.yonkersny.gov/Home/Components/News/News/5284/3232

ACKNOWLEDGMENTS. We thank J. Mueller, D. Politopoulos, C. Walker, and L. Yakaboski at Yonkers Police Department for their help fielding the officer surveys and their comments and feedback during the research process. We thank L. Bassan-Nygate, S. Riaz, C. Federico for helpful comments and feedback. This research was presented and benefited from feedback at the Institute for Humanities and Social Sciences and the 79th Midwest Political Science Association Annual Meeting. K.P. and P.V. received funding and research support for the studies in Yonkers, NY from the Justice Collaboratory (JC) at Yale Law School, where they previously held post-doctoral research fellowships (2020-2021). These studies were embedded in a panel survey initiative started by the JC in 2020, which was reviewed by the Human Subjects Committee at Yale University (IRB Protocol ID 2000028778) and determined to be exempt under federal regulation 45 CFR 46.104(d)(2)(iii). Consent to participate in panel surveys was obtained from participating police officers and community residents during the baseline survey wave. The study conducted on the national sample was embedded in a cooperative survey funded by the Center for the Study of American Politics and the Institution for Social and Policy Studies at Yale University (IRB Protocol ID 1312013102). All replication data and code are available at K.P.'s Dataverse within the Harvard Dataverse Network: https://dataverse.harvard.edu/dataverse/kpeyton.

1. PR Center, Behind the badge: Amid protests and calls for reform, how police view their jobs, key issues and recent fatal encounters between blacks and police. *Pew Res.* (2017).
2. K Peyton, M Sierra-Arevalo, DG Rand, A field experiment on community policing and police legitimacy. *Proc. Natl. Acad. Sci.* **116**, 19894–19898 (2019).
3. BA Ba, D Knox, J Mummolo, R Rivera, The role of officer race and gender in police-civilian interactions in chicago. *Science* **371**, 696–702 (2021).
4. PE Vaughn, K Peyton, HG A., Mass support for proposals to reshape policing depends on the implications for crime and safety. *Criminol. & Public Policy* **21**, 125–146 (2022).
5. S Karim, Relational state building in areas of limited statehood: Experimental evidence on the attitudes of the police. *Am. Polit. Sci. Rev.* **114**, 536–551 (2020).
6. D Yokum, A Ravishankar, A Coppock, A randomized control trial evaluating the effects of police body-worn cameras. *Proc. Natl. Acad. Sci.* **116**, 10329–10332 (2019).
7. MC Williams Jr, et al., Body-worn cameras in policing: Benefits and costs. *NBER Work. Pap.* (2021).
8. G Wood, TR Tyler, AV Papachristos, Procedural justice training reduces police use of force and complaints against officers. *Proc. Natl. Acad. Sci.* **117**, 9815–9821 (2020).
9. DS Nagin, CW Telep, Procedural justice and legal compliance: A revisionist perspective. *Criminol. & Public Policy* **19**, 761–786 (2020).
10. UD of Justice Equal Employment Opportunity Commission, , et al., Advancing diversity in law enforcement (2016).
11. NM Riccucci, GG Van Ryzin, H Li, Representative bureaucracy and the willingness to coproduce: An experimental study. *Public Adm. Rev.* **76**, 121–130 (2016).
12. NM Riccucci, GG Van Ryzin, K Jackson, Representative bureaucracy, race, and policing: A survey experiment. *J. Public Adm. Res. Theory* **28**, 506–518 (2018).
13. AR Miller, C Segal, Do female officers improve law enforcement quality? effects on crime reporting and domestic violence. *The Rev. Econ. Stud.* **86**, 2220–2247 (2019).
14. SS Hyland, E Davis, Local police departments, 2016: Personnel. *Las Vegas: United States Dep. Justice Off. Justice Programs Bureau Justice Stat.* (2019).
15. PTF on 21st Century Policing, Final report of the president's task force on 21st century policing (2015).
16. L Leatherby, R Oppel, Which police departments are as diverse as their communities. *New York Times* (2020).
17. B Ba, et al., Who are the police? descriptive representation in the coercive arm of government. *Work. Pap.* (2022).
18. JT Pickett, Public opinion and criminal justice policy: Theory and research. *Annu. Rev. Criminol.* **2**, 405–428 (2019).
19. JO Brehm, S Gates, *Working, shirking, and sabotage: Bureaucratic response to a democratic public.* (University of Michigan Press), (1999).
20. DP Rosenbaum, WP McCarty, Organizational justice and officer "buy in" in american policing. *Policing: an international journal police strategies & management* **40**, 71–85 (2017).
21. MW Kraus, JM Rucker, JA Richeson, Americans misperceive racial economic equality. *Proc. Natl. Acad. Sci.* **114**, 10324–10331 (2017).
22. MW Kraus, IN Onyeador, NM Daumeyer, JM Rucker, JA Richeson, The misperception of racial economic inequality. *Perspectives on Psychol. Sci.* **14**, 899–921 (2019).
23. IN Onyeador, et al., Disrupting beliefs in racial progress: Reminders of persistent racism alter perceptions of past, but not current, racial economic equality. *Pers. Soc. Psychol. Bull.* **47**, 753–765 (2021).
24. C Wong, J Bowers, T Williams, KD Simmons, Bringing the person back in: Boundaries, perceptions, and the measurement of racial context. *The J. Polit.* **74**, 1153–1170 (2012).
25. S Davidai, J Walker, Americans misperceive racial disparities in economic mobility. *Pers. Soc. Psychol. Bull.* p. 01461672211024115 (2021).
26. D Knox, W Lowe, J Mummolo, Administrative records mask racially biased policing. *Am. Polit. Sci. Rev.* **114**, 619–637 (2020).
27. CA Gallagher, Miscounting race: Explaining whites' misperceptions of racial group size. *Sociol. Perspectives* **46**, 381–396 (2003).
28. R Kardosh, AY Sklar, A Goldstein, Y Pertzov, RR Hassin, Minority salience and the overestimation of individuals from minority groups in perception and memory. *Proc. Natl. Acad. Sci.* **119**, e2116884119 (2022).

29. S Ansolabehere, M Meredith, E Snowberg, Asking about numbers: Why and how. *Polit. Analysis* **21**, 48–69 (2013).
30. JH Kuklinski, PJ Quirk, J Jerit, D Schwieder, RF Rich, Misinformation and the currency of democratic citizenship. *The J. Polit.* **62**, 790–816 (2000).
31. JG Bullock, Elite influence on public opinion in an informed electorate. *Am. Polit. Sci. Rev.* **105**, 496–515 (2011).
32. I Haaland, C Roth, J Wohlfart, Designing information provision experiments. *J. Econ. Lit.* (2021).
33. A Grigorieff, C Roth, D Ubfal, Does information change attitudes toward immigrants? *Demography* **57**, 1117–1143 (2020).
34. I Haaland, C Roth, Labor market concerns and support for immigration. *J. Public Econ.* **191**, 104256 (2020).
35. I Haaland, C Roth, Beliefs about racial discrimination and support for pro-black policies. *The Rev. Econ. Stat.* pp. 1–38 (2021).
36. B Callaghan, L Harouni, CH Dupree, MW Kraus, JA Richeson, Testing the efficacy of three informational interventions for reducing misperceptions of the black–white wealth gap. *Proc. Natl. Acad. Sci.* **118** (2021).
37. M Vlasceanu, MJ Morais, A Coman, The effect of prediction error on belief update across the political spectrum. *Psychol. Sci.* **32**, 916–933 (2021).
38. NM Riccucci, GG Van Ryzin, CF Lavena, Representative bureaucracy in policing: Does it increase perceived legitimacy? *J. public administration research theory* **24**, 537–551 (2014).
39. GG Van Ryzin, NM Riccucci, H Li, Representative bureaucracy and its symbolic effect on citizens: a conceptual replication. *Public Manag. Rev.* **19**, 1365–1379 (2017).
40. M Sievert, A replication of "representative bureaucracy and the willingness to coproduce". *Public Adm.* **99**, 616–632 (2021).
41. K Peyton, Does trust in government increase support for redistribution? evidence from randomized survey experiments. *Am. Polit. Sci. Rev.* **114**, 596–602 (2020).
42. MA Craig, JA Richeson, On the precipice of a "majority-minority" america: Perceived status threat from the racial demographic shift affects white americans' political ideology. *Psychol. science* **25**, 1189–1197 (2014).
43. TL Dover, B Major, CR Kaiser, Members of high-status groups are threatened by pro-diversity organizational messages. *J. Exp. Soc. Psychol.* **62**, 58–67 (2016).
44. ND Brown, DS Jacoby-Senghor, Majority members misperceive even "win-win" diversity policies as unbeneficial to them. *J. personality social psychology* (2021).
45. J Sunshine, TR Tyler, The role of procedural justice and legitimacy in shaping public support for policing. *Law & society review* **37**, 513–548 (2003).
46. JM Jones, In us, black confidence in police recovers from 2020 low. *Gallup. com* (2021).
47. GA Huber, C Paris, Assessing the programmatic equivalence assumption in question wording experiments: Understanding why americans like assistance to the poor more than welfare. *Public Opin. Q.* **77**, 385–397 (2013).
48. FJ Crosby, A Iyer, S Sincharoen, Understanding affirmative action. *Annu. Rev. Psychol.* **57**, 585–611 (2006).
49. DA Kravitz, SL Klineberg, Reactions to two versions of affirmative action action among whites, blacks, and hispanics. *J. Appl. Psychol.* **85**, 597 (2000).
50. H Haley, J Sidanius, The positive and negative framing of affirmative action: A group dominance perspective. *Pers. Soc. Psychol. Bull.* **32**, 656–668 (2006).
51. M Bertrand, E Duflo, Field experiments on discrimination. *Handb. economic field experiments* **1**, 309–393 (2017).
52. J Hainmueller, DJ Hopkins, The hidden american immigration consensus: A conjoint analysis of attitudes toward immigrants. *Am. J. Polit. Sci.* **59**, 529–548 (2015).
53. J Hemker, A Rink, Multiple dimensions of bureaucratic discrimination: Evidence from german welfare offices. *Am. J. Polit. Sci.* **61**, 786–803 (2017).
54. K Bansak, J Hainmueller, DJ Hopkins, T Yamamoto, Conjoint survey experiments. *Adv. experimental political science* **19** (2021).
55. DL Teele, J Kalla, F Rosenbluth, The ties that double bind: social roles and women's underrepresentation in politics. *Am. Polit. Sci. Rev.* **112**, 525–541 (2018).
56. S Schwarz, A Coppock, What have we learned about gender from candidate choice experiments? a meta-analysis of sixty-seven factorial survey experiments. *J. Polit.* (2022).
57. J Hainmueller, D Hangartner, T Yamamoto, Validating vignette and conjoint survey experiments against real-world behavior. *Proc. Natl. Acad. Sci.* **112**, 2395–2400 (2015).
58. Y Horiuchi, Z Markovich, T Yamamoto, Does conjoint analysis mitigate social desirability bias? *Polit. Analysis* pp. 1–15 (2021).
59. G Blair, A Coppock, M Moor, When to worry about sensitivity bias: A social reference theory and evidence from 30 years of list experiments. *Am. Polit. Sci. Rev.* **114**, 1297–1315 (2020).
60. MD White, JA Cooper, J Saunders, AJ Raganella, Motivations for becoming a police officer: Re-assessing officer attitudes and job satisfaction after six years on the street. *J. criminal justice* **38**, 520–530 (2010).
61. JC Gibbs, Diversifying the police applicant pool: Motivations of women and minority candidates seeking police employment. *Crim. Justice Stud.* **32**, 207–221 (2019).
62. J Hainmueller, DJ Hopkins, T Yamamoto, Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Polit. Analysis* **22**, 1–30 (2014).

Supporting Information for Online Appendix

Kyle Peyton* Chagai Weiss † Paige Vaughn ‡

Version: April 2022

Contents

*Research Fellow in Political Science, Institute for Humanities and Social Sciences, Australian Catholic University, Email: kyle.peyton@acu.edu.au

†PhD candidate in Political Science, University of Wisconsin – Madison; Middle East Initiative Predoctoral Fellow, Harvard Kennedy School. Email: cmweiss@wisc.edu

‡Assistant Professor of Sociology and Criminology, Spring Hill College. Email: pvaughn@shc.edu

List of Figures

List of Tables

S1 Survey methodology and experimental designs

Our results are based on data collected from a series of surveys with embedded experiments. These include one survey on a national sample of the U.S. adult population, two municipal surveys fielded on a sample of the adult population of Yonkers, NY, and one survey fielded on officers from the Yonkers Police Department (YPD). Section S1.1 provides a description of the survey methodology for the national experiment and Section S1.2 provides a description of the survey methodology for municipal surveys.

The information provision experiments that were embedded in the national survey and the second municipal survey of Yonkers residents are described in Section S1.3, along with the (pre-treatment) questions used to measure respondents' prior beliefs about police diversity and our (post-treatment) outcome measures. The conjoint experiments on police officer recruitment that were embedded in the first municipal survey of Yonkers residents and the survey of YPD officers are described in Section S1.4.

S1.1 National survey

Survey data were collected in July 2021 using the Lucid platform. Lucid is an aggregator of survey respondents that sources from a wide variety of providers. Participants are paid directly by the vendor, either in US dollars or through a points program at a similar rate. Lucid provides customers with quota samples that approximate US census margins on age, gender, race/ethnicity and region, and these samples compare favorably with nationally representative samples on demographic, psychological, and experimental estimates [1]. Importantly, experimental estimates obtained from online quota samples and other non-probability samples closely match those obtained from probability samples, demonstrating that causal effects estimated via randomized experiments on these samples typically generalize to the broader population of interest [2, 3, 4, 5, 1, 6].

Following best practices to ensure data quality in online survey sampling [7, 8, 6, 9], we restricted participation to respondents that passed an attention screener placed at the beginning of the survey. Screening out respondents based on attention checks placed near the end of the survey or after treatments are administered can induce bias, but using attention checks administered early in the survey or prior to experimental treatments does not induce bias [10, 11]. We employed an attention screener used in recent survey experimental work on the Lucid platform [6, 12]. After viewing the screener (shown in Fig. S1), participants were asked the following two attention check questions (the correct responses are highlighted in bold text):

1. How was Simon identified by police for the crime he allegedly committed? [A police officer recognized him, From video surveillance, **Because he left his ID**, He turned himself in, None of the above]
2. How much money did Simon allegedly steal? [About \$500, **About \$1500**, About \$25,000, About \$1 million dollars, None of the above]

Among the 3,373 individuals that consented to participate in the survey, 60% passed the first attention check question (ACQ) and went on to complete our survey (this was a cooperative survey with another research team, and the first half of the survey was allocated to demographic questions and this team's survey content). Among these 2,017 individuals, 60% ($n = 1,219$) also passed the second ACQ. Only those individuals that provided incorrect answers to the first ACQ, or refused to answer, were terminated from the survey. Among these 2,017 individuals, the median time to completion was 14 minutes.

We constructed survey weights to adjust for potential differences in respondent demographics along the following characteristics: sex, region, Hispanic, race, household income, education, and age. Our target proportions for these characteristics are the estimates reported by the American Community Survey (ACS) and the U.S. Census. Weights are constructed using the autumn package in R [13], which implements an iterative raking procedure used by the American National Election Study (ANES) survey [14].

Applying the weights reduced the average difference between the sample proportions and target marginals from 0.02 to 0.003 for an effective sample size of 1,437 units from a nominal size of 2,017 (implying a design effect of 1.40). Table S1 compares the unweighted sample proportions, weighted sample proportions, and the target proportions across background covariates. Given that our focus is on estimating causal effects, we do not use survey weights for any analyses presented here or in the manuscript [15].

MAN ARRESTED FOR STRING OF BANK THEFTS

Columbus Police have arrested a man they say gave his driver's license to a teller at a bank he was robbing.

According to court documents, Bryan Simon is accused of robbing four Central Ohio banks between October 3 and November 5, 2018.

During a robbery on November 5 at the Huntington Bank, the sheriff's office says Simon was tricked into giving the teller his drivers' license.

According to court documents, Simon approached the counter and presented a demand note for money that said "I have a gun." The teller gave Simon about $500, which he took.

Documents say Simon then told the teller he wanted more money. The teller told him a driver's license was required to use the machine to get our more cash. Simon reportedly then gave the teller his license to swipe through the machine and then left the bank with about $1000 in additional cash, but without his ID.

Detectives arrested him later that day at the address listed on his ID.

Figure S1: Attention screener used in Lucid sample.

	Sample proportion	**Target proportion**	**Absolute deviation**
Sex:			
Female	0.52	0.52	0.00
Male	0.48	0.48	0.00
Race/ethnicity:			
White	0.65	0.64	0.01
Hispanic	0.16	0.17	0.01
Black	0.12	0.12	0.00
AAPI	0.03	0.06	0.03
Other	0.03	0.01	0.02
Age:			
18-24	0.13	0.12	0.01
25-29	0.09	0.10	0.01
30-34	0.10	0.09	0.01
35-39	0.08	0.09	0.01
40-44	0.08	0.08	0.00
45-49	0.08	0.08	0.00
50-54	0.05	0.08	0.03
55-59	0.07	0.08	0.01
60-64	0.10	0.08	0.02
65-69	0.07	0.07	0.00
70-74	0.08	0.06	0.02
75+	0.06	0.08	0.02
Region:			
South	0.40	0.38	0.02
West	0.24	0.24	0.00
Midwest	0.19	0.21	0.02
Northeast	0.17	0.18	0.00
Educational attainment:			
No high school diploma	0.07	0.10	0.03
High school diploma	0.42	0.45	0.03
Associate's degree	0.09	0.10	0.01
Bachelor's degree	0.22	0.22	0.00
Graduate degree	0.19	0.13	0.06
Income:			
$15,000 or less	0.18	0.09	0.09
$15,000-$24,999	0.10	0.09	0.01
$25,000-$34,999	0.11	0.08	0.03
$35,000-$49,999	0.10	0.12	0.02
$50,000-$74,999	0.15	0.17	0.02
$75,000-$99,999	0.12	0.12	0.00
$100,000-$149,999	0.11	0.15	0.04
$150,000-$199,999	0.06	0.08	0.03
$200,000 and above	0.02	0.10	0.08

Table S1: Demographic characteristics for the national sample and the U.S. adult population. The target proportions for the U.S. adult population (18+) come from the American Community Survey and the U.S. Census.

S1.2 Municipal surveys

Our survey data on Yonkers residents were collecting as part of the "Community Vitality Survey" initiative started by Yale Law School in 2021. The broad aim of this survey initiative was to conduct municipal surveys of police officers and the residents they police across different U.S. cities to better understand public opinion on policing at the local level. Although the survey initiative was primarily aimed at collecting descriptive data on the views of residents and police officers, several experiments were also embedded in the early survey waves (which form the basis of our analyses).

Survey data for the resident population were collected using the mail to online panel design, which uses public information – sourced primarily from voter registration records – to construct a baseline sample frame of the adult population [16, 17, 18, 19]. We started by recruiting participants in Yonkers, NY to participate in an online panel survey called the "Yonkers Community Vitality Survey". Between May and July 2021, 63,743 residents were sent an invitation to the mailing address listed in a voter-file purchased from L2 political, a commercial vendor. The mailers directed recipients to an online survey via a landing page at a dedicated university website.[1] Each respondent was provided with a unique login code in their recruitment letter, and a dedicated phone number and university email address were created to field respondent inquiries during the recruitment period. 1,413 individuals completed this initial baseline survey, for a response rate of approximately 2.2%. Response rates of between 2-3% are common in previous studies that have used the mail to online panel design[16, 18], though response rates above 5% have also been achieved [18].

We constructed survey weights to adjust for potential differences in respondent demographics along the following characteristics: sex, race/ethnicity, age, birthplace, education, and income. Our target proportions for these characteristics are the estimates reported by the American Community Survey (ACS) and the U.S. Census. Weights are constructed using the autumn package in R [13], which implements an iterative raking procedure used by the American National Election Study (ANES) survey [14]. Applying the weights reduced the average difference between the sample proportions and target marginals from 0.05 to 0.01 for an effective sample size of 594 units from a nominal size of 1,413 (implying a design effect of 2.38). Table S2 compares the unweighted sample proportions, weighted sample proportions, and the target proportions across background covariates. Given that our focus is on estimating causal effects, we do not use survey weights for any analyses presented here or in the manuscript [15].

In October 2021, approximately three months after the baseline survey responses were collected, all individuals that responded to the baseline survey were invited via email to complete a follow-up survey. 644 of the 1,413 individuals that completed the baseline survey also completed the follow-up survey, for a return response rate of 46%. The police recruitment conjoint experiment (detailed in Section S1.4) was embedded in the first survey, and the information provision experiment (detailed in Section S1.3) was embedded in the follow-up survey. Table S3 compares the available demographic characteristics of the officer population with the survey sample.

Survey data for the officer population were collected via invitations delivered directly to the

[1] www.communityvitality.yale.edu

government email addresses of 600 sworn police officers, which were provided to the researchers by the police department. Of the 600 officers invited to participate in the survey, 250 completed the survey for a response rate of 42%. For context, response rates in web-based surveys of police officers averaged about 40% between 1996 and 2016, and these rates have been declining over time [20]. The police recruitment conjoint experiment (detailed in Section S1.4) was embedded in this survey.

All surveys were administered using Qualtrics Survey Software, and respondents could choose to either receive a fixed payment of $5 or enter a raffle to win one of 13 $70 payments. All payments were delivered via TangoCard, a commercial vendor of electronic gift cards.

	Sample proportion	Target proportion	Absolute deviation
Sex:			
Female	0.59	0.53	0.06
Male	0.41	0.47	0.06
Race/ethnicity:			
White	0.45	0.36	0.10
Hispanic	0.29	0.39	0.09
Black	0.14	0.16	0.01
AAPI	0.06	0.07	0.00
Other	0.05	0.03	0.02
Age:			
18-24	0.10	0.12	0.03
25-29	0.08	0.09	0.01
30-34	0.11	0.10	0.02
35-44	0.16	0.17	0.01
45-54	0.16	0.17	0.01
55-64	0.16	0.15	0.01
65-74	0.15	0.10	0.05
75+	0.08	0.09	0.01
Birthplace:			
United States	0.77	0.62	0.16
Another country	0.23	0.38	0.16
Educational attainment:			
No high school diploma	0.02	0.17	0.15
High school diploma	0.28	0.44	0.17
Associate's degree	0.08	0.08	0.00
Bachelor's degree	0.32	0.18	0.13
Graduate degree	0.31	0.13	0.18
Income:			
$15,000 or less	0.14	0.13	0.00
$15,000-$24,999	0.08	0.09	0.01
$25,000-$34,999	0.09	0.08	0.01
$35,000-$49,999	0.13	0.11	0.02
$50,000-$74,999	0.20	0.15	0.05
$75,000-$99,999	0.15	0.12	0.04
$100,000-$149,999	0.14	0.16	0.01
$150,000-$199,999	0.04	0.08	0.04
$200,000 and above	0.03	0.09	0.06

Table S2: Demographic characteristics for municipal sample and adult population in Yonkers, NY. The target proportions for the Yonkers adult population (18+) come from the American Community Survey and the U.S. Census.

	Sample proportion	Target proportion	Absolute deviation
Sex:			
Female	0.15	0.14	0.01
Male	0.85	0.86	0.01
Race/ethnicity:			
White	0.82	0.78	0.04
Hispanic	0.12	0.15	0.03
Black	0.06	0.07	0.01
AAPI	-	<0.01	<0.01
Age:			
18-24	0.02	0.02	0.00
25-29	0.10	0.10	0.00
30-34	0.19	0.20	0.01
35-44	0.42	0.39	0.03
45-54	0.24	0.25	0.01
55-64	0.03	0.04	0.01

Table S3: Demographic characteristics for police officer sample and the population of police officers employed at YPD.

S1.3 Information provision experiments

The information provision experiment fielded in the national survey randomly assigned 2,017 individuals to one of four possible conditions: 1) no information (control); 2) information about police diversity only (hereafter "Info treatment"); 3) information about a recent *Science* publication [21] describing the potential benefits of police diversification for minority residents ("*Science* treatment"); 4) both information about police diversity and the *Science* article ("Info + *Science*"). Prior to treatment assignment, respondents beliefs about police officer diversity were measured using the questions presented in Fig. S2. Those assigned to the no information (control) condition simply reported their beliefs and did not receive any additional information. Fig. S3 shows the information treatment and Fig. S4 shows the *Science* treatment. Those assigned to the "Info + *Science*" condition received both, presented in randomized order. For comparison with the information experiment fielded in the municipal sample (which only involved two conditions), we restrict attention to effects of the Info condition (relative to control) in the manuscript. We provide a complete analysis of all treatment effects in Section S2.1.5. This experiment was not pre-registered.

The information provision experiment fielded in the municipal survey randomly assigned 644 individuals to one of two possible conditions: 1) no information (control); or 2) information about police officer diversity at their local police department (information treatment). Prior to treatment assignment, respondents beliefs about police officer diversity at their local police department were measured using the questions presented in Fig. S6. Those assigned to the no information (control) condition simply reported their beliefs and did not receive any additional information. Those assigned to the information condition received accurate information about police officer diversity, alongside the estimates they provided (see Fig. S6). In anticipation of sample size constraints, we did not include the additional two treatment arms from the information experiment fielded on the national sample. This experiment was pre-registered (see Section S3 for pre-registration).

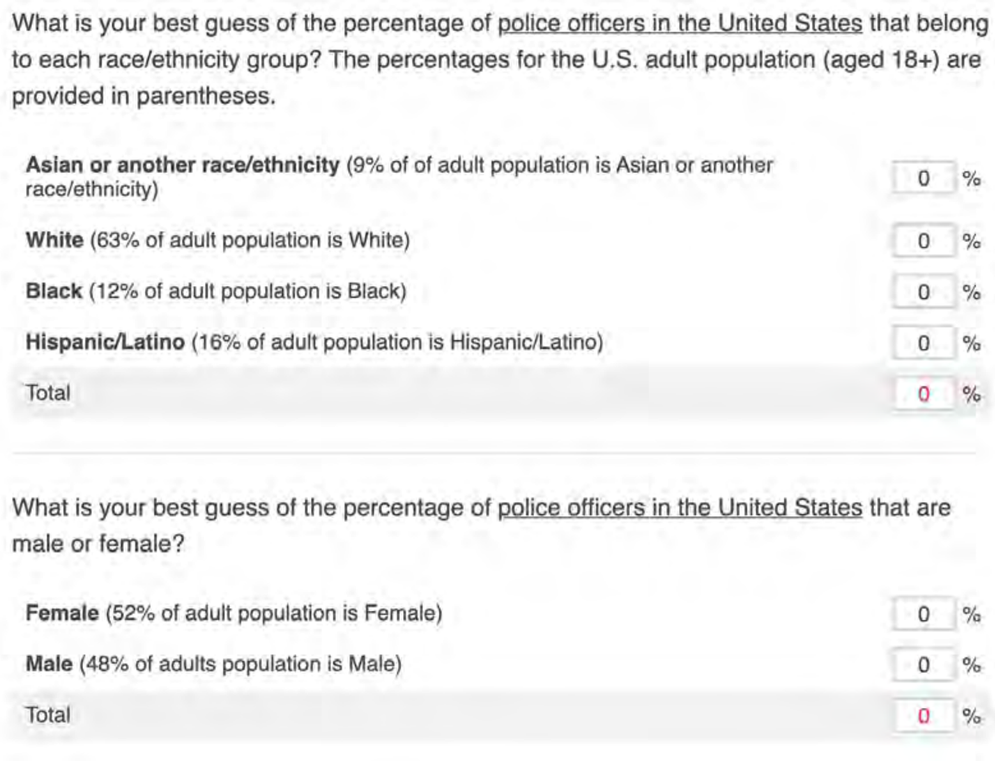

What is your best guess of the percentage of police officers in the United States that belong to each race/ethnicity group? The percentages for the U.S. adult population (aged 18+) are provided in parentheses.

Asian or another race/ethnicity (9% of of adult population is Asian or another race/ethnicity)	0	%
White (63% of adult population is White)	0	%
Black (12% of adult population is Black)	0	%
Hispanic/Latino (16% of adult population is Hispanic/Latino)	0	%
Total	0	%

What is your best guess of the percentage of police officers in the United States that are male or female?

Female (52% of adult population is Female)	0	%
Male (48% of adults population is Male)	0	%
Total	0	%

Figure S2: Pre-treatment measures of beliefs about police officer diversity in national sample

The following table shows how the estimates for U.S. police officer demographics that you provided at the beginning of the survey compare to the most recent data published by the U.S. Department of Justice. The percentages for the U.S. adult population (aged 18+) are again provided as a reference.

	Your estimate	**U.S. police**	**U.S. adults**
Race/ethnicity:			
White	63%	72%	63%
Black	12%	11%	12%
Hispanic/Latino	16%	13%	16%
Asian or other	9%	4%	9%
Sex:			
Male	70%	88%	48%
Female	30%	12%	52%

Figure S3: Example treatment assignment for information condition in national sample. Values in the "your estimate" column are provided for illustrative purposes and correspond to the sample medians.

A team of researchers from Princeton, Columbia, and University of California-Irvine recently conducted a study to examine whether the race and gender of officers and civilians affect their interactions. The study found that, relative to White officers, Black and Hispanic officers working in similar conditions made fewer stops and arrests, and used force less often, especially against Black civilians. The study also found that female officers used less force than male officers across all racial groups.

The study was published in February 2021 at *Science*, one of the world's top research outlets. A summary is provided below.

RESEARCH ARTICLE

CRIMINAL JUSTICE

The role of officer race and gender in police-civilian interactions in Chicago

Bocar A. Ba[1], Dean Knox[2]*, Jonathan Mummolo[3]*, Roman Rivera[4]

Diversification is a widely proposed policing reform, but its impact is difficult to assess. We used records of millions of daily patrol assignments, determined through fixed rules and preassigned rotations that mitigate self-selection, to compare the average behavior of officers of different demographic profiles working in comparable conditions. Relative to white officers, Black and Hispanic officers make far fewer stops and arrests, and they use force less often, especially against Black civilians. These effects are largest in majority-Black areas of Chicago and stem from reduced focus on enforcing low-level offenses, with greatest impact on Black civilians. Female officers also use less force than males, a result that holds within all racial groups. These results suggest that diversity reforms can improve police treatment of minority communities.

Figure S4: Screenshot of *Science* treatment arm in national sample. Those assigned to the "Info + *Science*" treatment arm received this and the information treatment from Fig. S3 in randomized order.

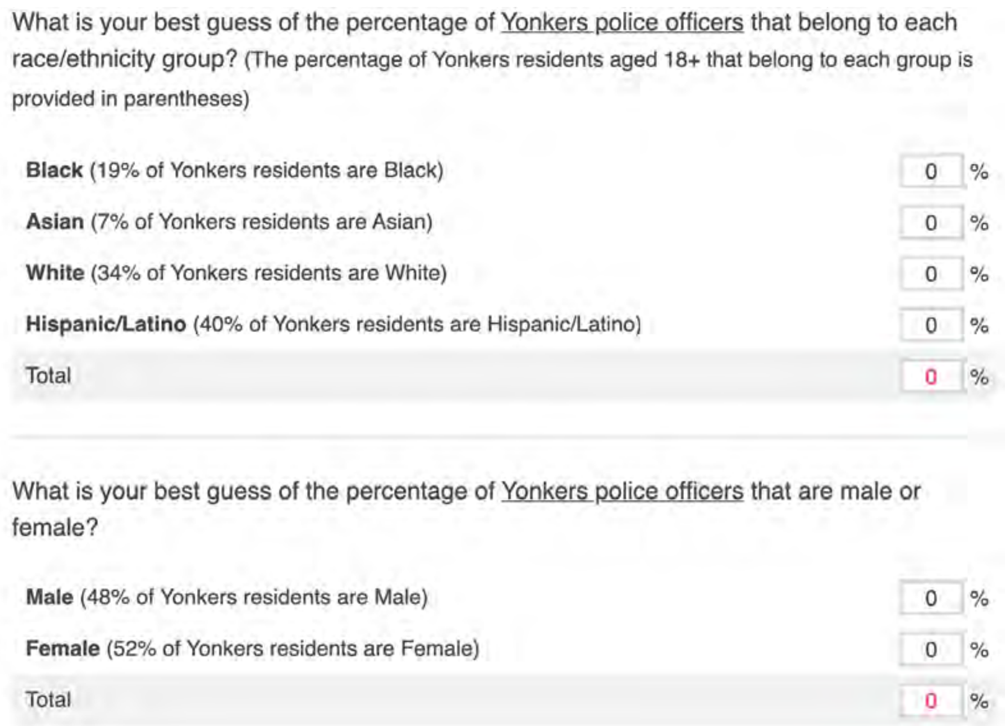

What is your best guess of the percentage of Yonkers police officers that belong to each race/ethnicity group? (The percentage of Yonkers residents aged 18+ that belong to each group is provided in parentheses)

Black (19% of Yonkers residents are Black)	0	%
Asian (7% of Yonkers residents are Asian)	0	%
White (34% of Yonkers residents are White)	0	%
Hispanic/Latino (40% of Yonkers residents are Hispanic/Latino)	0	%
Total	0	%

What is your best guess of the percentage of Yonkers police officers that are male or female?

Male (48% of Yonkers residents are Male)	0	%
Female (52% of Yonkers residents are Female)	0	%
Total	0	%

Figure S5: Pre-treatment measures of beliefs about police officer diversity in municipal sample

We will now provide you with some information about diversity in the Yonkers Police Department. The table below shows how the estimates for police officer demographics that you provided earlier in the survey compare to the true percentages. The percentages for the entire Yonkers community (adult population aged 18+) are again provided as a reference.

	Your estimate	Yonkers Police	Yonkers Community
Race/ethnicity:			
White	60%	78%	34%
Hispanic/Latino	20%	15%	40%
Black	15%	6%	19%
Asian or other	5%	<1%	7%
Sex:			
Male	75%	86%	48%
Female	25%	14%	52%

YPD employs 626 full-time officers: 490 White, 95 Hispanic/Latino, 38 Black, 3 Asian I 538 Male, 88 Female

Figure S6: Example treatment assignment for information condition in municipal sample. Values in the "your estimate" column are provided for illustrative purposes and correspond to the sample medians.

S1.3.1 Behavioral outcomes and survey indexes in manuscript

Our two behavioral outcome measures – voting for police department diversification and donations to Black police officer association – appeared near the end of the survey, after the survey-based outcome measures. The first is a binary outcome coded 1 if respondents chose police department diversification from a list of four potential policy changes and 0 otherwise (see Fig. S7). The second behavioral outcome is the dollar amount of a potential bonus payment (between $0 and $50) that the respondent would donate (as opposed to keep) to a local non-profit supporting Black police officers (see Fig. S8).

The four outcome indexes reported in the manuscript were constructed from multiple survey items used in prior work on attitudes towards police [19], and each was combined into a single index using inverse covariance weighting [22]. The question wordings for each individual item are provided below. The first two outcome indexes were measured in both the municipal and national samples. The second two outcome indexes were measured in the municipal sample, in both the baseline and followup survey.

1. **Support for affirmative action in recruitment and hiring (4-item index)**. Prompt: "To what extent do you support or oppose implementing affirmative action programs to increase recruitment and hiring of police officers at YPD from each of the following groups".

Respondents then reported their support for four groups: “Female officers”, “Black officers”, “Hispanic/Latino officers”, and “Asian officers”. Each was presented in random order and support was recorded using a 7 point scale from “Strongly disagree” to “Strongly agree” with a neutral midpoint. $\alpha = 0.98$ in municipal sample and $\alpha = 0.96$ in national sample.

2. **Support for tie-breaking in favor of minority applicants (4-item index)**. Each item was presented in random order. Prompt: “Imagine YPD is trying to decide between two equally qualified applicants for police officer. For each of the comparisons listed below, please say what you think YPD should do.” $\alpha = 0.89$ in municipal sample and $\alpha = 0.81$ in national sample.

    - “Two equally qualified applicants: one female and the other male. ” [1 = “Hire the female applicant”, -1 = “Hire the male applicant”, 0 = “Random selection (e.g., let a coin flip decide)”]
    - “Two equally qualified applicants: one Black and the other White.” [1 = “Hire the Black applicant”, -1 = “Hire the White applicant”, 0 = “Random selection (e.g., let a coin flip decide)”]
    - “Two equally qualified applicants: one Hispanic/Latino and the other White.” [1 = “Hire the Hispanic/Latino applicant”, -1 = “Hire the White applicant”, 0 = “Random selection (e.g., let a coin flip decide)”]
    - “Two equally qualified applicants: one Asian and the other White.” [1 = “Hire the Asian applicant”, -1 = “Hire the White applicant”, 0 = “Random selection (e.g., let a coin flip decide)”]

3. **Trust and confidence in the local police department (2-item index)**. Each item below was presented in random order, with responses recorded using the 5-point scales in brackets. $\alpha = 0.80$ in baseline survey and $\alpha = 0.83$ in followup survey.

    - “How much of the time do you think Yonkers residents can trust the Yonkers Police Department to do what is right?” [1 = “Never”, 2 = “Sometimes”, 3 = “About half the time”, 4 = “Most of the time”, 5 = “Always”]
    - “How much confidence do you have in Yonkers Police Department to act in the best interest of the public?” [1 = “None”, 2 = “Very little”, 3 = “Some”, 4 = “Quite a lot”, 5 = “A great deal”]

4. **Willingness to cooperate with police (4-item index)**. Each of the four items below were presented in random order and responses were recorded using a 7 point scale from “Extremely unlikely” to “Extremely likely” with a neutral midpoint. $\alpha = 0.74$ in baseline survey and $\alpha = 0.73$ in followup survey.

    - “If the police were looking for a suspect who was hiding, and you knew where that person was, how likely would you be to provide the police with information?”
    - “How likely would you be to call the police to report a crime?”
    - “How likely would you be to report suspicious activity to the police?”
    - “How likely would you be to attend a community meeting to discuss problems in your neighborhood with the police?”

We will now provide you with an opportunity to express your views about police reform to your local representatives in Yonkers.

Please select from the list below what you would prefer to see prioritized at YPD. If you would like to see something else prioritized, please select "Something else" and type in a short description. If you don't want to participate, select "don't want to participate".

After this survey is complete, we will tally the results from all participants and present these data directly to Mayor Mike Spano and Police Commissioner John Mueller.

Note: please remember that your responses are anonymous, and your identity is protected. These data will be presented in aggregate form, so it is not possible to identify individual participants.

- ○ **Civilian oversight**. Create a Civilian Review Board with the power to investigate and recommend action for complaints made against police officers for instances that include excessive force, abuse of authority, and offensive language.
- ○ **Diversification**. Implement affirmative action programs that increase recruitment and hiring of officers from underrepresented groups so that police more closely resemble the community in terms of race/ethnicity and gender.
- ○ **Community policing**. Establish regular meetings between police and the public that provide a forum for city representatives, businesses, and residents to share information and cooperatively address neighborhood issues.
- ○ **Body worn cameras**. Require police officers to wear body cameras that record their interactions with the public while on duty.
- ○ Something else not listed
- ○ Don't want to participate

Figure S7: Voting for police department diversification question

There are several non-profit organizations across the United States that work to increase the recruitment and retention of police officers from under-represented minority groups. One local organization, called the *Yonkers Guardians Association*, supports Black individuals working in law enforcement in Yonkers, NY.

One of the organization's key initiatives it to promote equal employment opportunities at the Yonkers Police Department through recruitment, appointments, assignments, and promotions. Below, you are given the opportunity to financially support the *Yonkers Guardians Association* through a donation.

Here's how it works: 1 out of every 20 people that complete this survey will be randomly selected to receive a payment of $50. If selected, you may keep the entire amount or donate any amount between $0 and $50 to the *Yonkers Guardians Association.* This is addition to what you will already be paid for completing this survey.

What amount (if any) would you like to donate to the *Yonkers Guardians Association*? (the total must sum to $50)

Amount to donate	$ 0
Amount to keep	$ 0
Total	$ 0

Figure S8: Donation to Black officers association question

S1.4 Police recruitment conjoint experiment

The police recruitment conjoint was embedded in the initial (baseline) municipal survey of 1,413 Yonkers residents, and a direct replication of the same experiment was then embedded in the survey of 250 police officers from the YPD (see Section S4 for pre-registration). In both surveys, respondents were first provided with a detailed description of the task they would be asked to complete and the information that would be provided to them (see Fig. S9).

For the next few minutes, we will provide you with several pieces of information about people who might apply to join the Yonkers PD. For each pair of people, please indicate which of the two applicants you would prefer to see recruited into the Yonkers PD.

Although this exercise is purely hypothetical, the information we provide is based on real police officer applications and therefore provides a realistic portrait of the types of people that might apply.

Please remember that police departments receive many more applications than they can accept. Even if you aren't entirely sure, please indicate which of the two applicants you prefer. **The information you will be provided is shown below.**

	Information that will be provided
Age	The age of each applicant at the time of their application
Sex	Whether the applicants are male or female
Race/Ethnicity	The race/ethnicity of each applicant
Education	The highest level of education each applicant had completed at the time of their application
Yonkers Resident	Whether the applicants currently live in Yonkers and, if so, for how long
Previous occupation	The occupation that each applicant held at the time of their application
Civil service exam	How each applicant performed on the civil service exam. For example, "Top 5 percent" means an applicant scored higher than 95% of all other applicants
Motivation for becoming a police officer	What each applicant stated as their reason for applying to become a police officer. These have been grouped into common categories to simplify comparisons

Figure S9: Task instructions for police recruitment conjoint

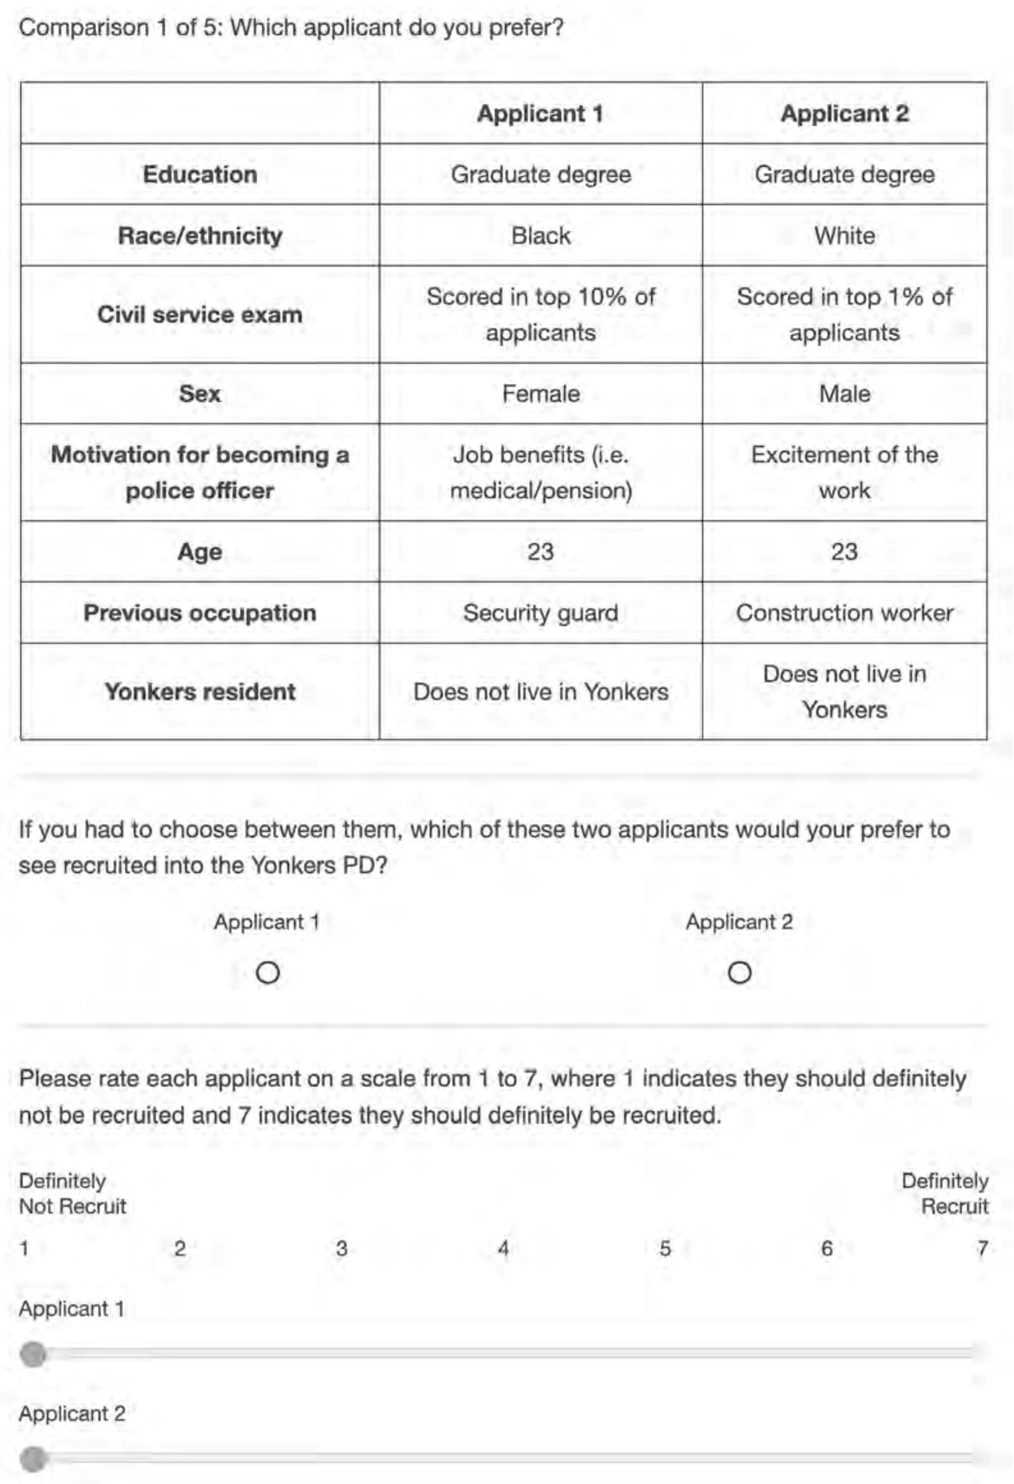
Comparison 1 of 5: Which applicant do you prefer?

	Applicant 1	Applicant 2
Education	Graduate degree	Graduate degree
Race/ethnicity	Black	White
Civil service exam	Scored in top 10% of applicants	Scored in top 1% of applicants
Sex	Female	Male
Motivation for becoming a police officer	Job benefits (i.e. medical/pension)	Excitement of the work
Age	23	23
Previous occupation	Security guard	Construction worker
Yonkers resident	Does not live in Yonkers	Does not live in Yonkers

If you had to choose between them, which of these two applicants would your prefer to see recruited into the Yonkers PD?

Applicant 1 ○ Applicant 2 ○

Please rate each applicant on a scale from 1 to 7, where 1 indicates they should definitely not be recruited and 7 indicates they should definitely be recruited.

Definitely Not Recruit — Definitely Recruit

1 2 3 4 5 6 7

Applicant 1

Applicant 2

Figure S10: Example round from police recruitment conjoint

Next, respondents evaluated five pairs of hypothetical police officer applicants, with the following randomly assigned features drawn for each attribute (in bold):

- **Age:** 23; 25; 27; 29; 31; 33; 35; 37
- **Sex**: Male; Female

- **Race/Ethnicity**: Asian; Black; Hispanic/Latino; White
- **Education**: GED; High school; Associates degree; Bachelors degree; Graduate degree
- **Yonkers Resident**: Does not live in city; For less than 1 year; For 1-2 years; For 3-5 years; For 6-10 years; For more than 10 years
- **Previous occupation**: Construction worker, Personal trainer, Server/Bartender, Retail salesperson, Security guard, Police officer in another city; Military service; School teacher; Social worker
- **Civil service exam**: Top 1% of applicants; Top 10% of applicants; Top 15% of applicants; Top 25% of applicants; Top 5% of applicants
- **Motivation for becoming a police officer**: Friends/relatives in police department; Excitement of the work; Lifelong dream/aspiration; To fight crime; Career advancement; Job benefits; Job security; To help people

Attributes were chosen based on a combination of interviews with the police officers at YPD recruitment division, historical data on real police officer applicants, and prior survey work on police officers' motivations and background characteristics [23, 24, 25]. In order to avoid implausible cases (e.g., school teacher's with GED's) we employed restricted randomization on the education and occupation attributes such that potential applicants that were previously school teachers or social workers always had education levels of at least a Bachelor's degree or higher. All estimates presented here and in the manuscript are adjusted to account for this conditional independence, which is a common feature in conjoint experiments that seek to avoid generating implausible profiles [26, 27]. Aside from this restriction, attributes were otherwise randomly assigned with uniform distribution.

S2 Supplementary analyses

S2.1 Information provision experiments

S2.1.1 Average treatment effects on survey index components

In the manuscript, we reported estimates of the Average Treatment Effect (ATE) on survey indexes used to measures support for affirmative action in hiring and recruitment, and support for tie-breaking in favor of minority applicants. As described in Section S1.3.1, these indexes were constructed using 4 separate question items that each focused on support for a specific under-represented minority group. Here, we conduct supplementary analyses that instead treat each index component as a unique item. Figure S11 shows these estimates (with 95% confidence intervals) and Table S4 reports the underlying point estimates and standard errors. To facilitate comparisons, all estimates are standardized using using Glass's Δ, which scales outcomes by the standard deviation in the control group [28, 29]. Although we find some numerical differences across the estimated ATEs on the individual components (e.g., larger point estimates for Black applicants) these are not statistically distinguishable from one another. This suggests that the precision gains we achieve from dimension reduction are worth the potential drawbacks associated with summary indexes that abstract away from effects on different groups.

	Municipal sample	**National sample**
Support for affirmative action in recruitment and hiring:		
Black officers	0.04 (0.08)	0.07 (0.06)
Hispanic/Latino officers	0.05 (0.08)	-0.02 (0.06)
Asian officers	0.02 (0.08)	-0.04 (0.06)
Female officers	0.01 (0.08)	-0.04 (0.06)
Support for tie-breaking in favor of minority applicants:		
Black applicants	0.26 (0.08)*	0.22 (0.06)*
Hispanic/Latino applicants	0.19 (0.08)*	0.14 (0.06)*
Asian applicants	0.24 (0.08)*	0.12 (0.07)
Female applicants	0.21 (0.08)*	0.06 (0.07)

Table S4: Estimated treatment effects on survey index components in each sample. Point estimates for the ATEs estimated using OLS regression of the outcome on treatment assignment, with robust standard errors in parentheses. All estimates are standardized using Glass's Δ, which scales outcomes by the standard deviation in the control group [28, 29].
$^*P < 0.05$

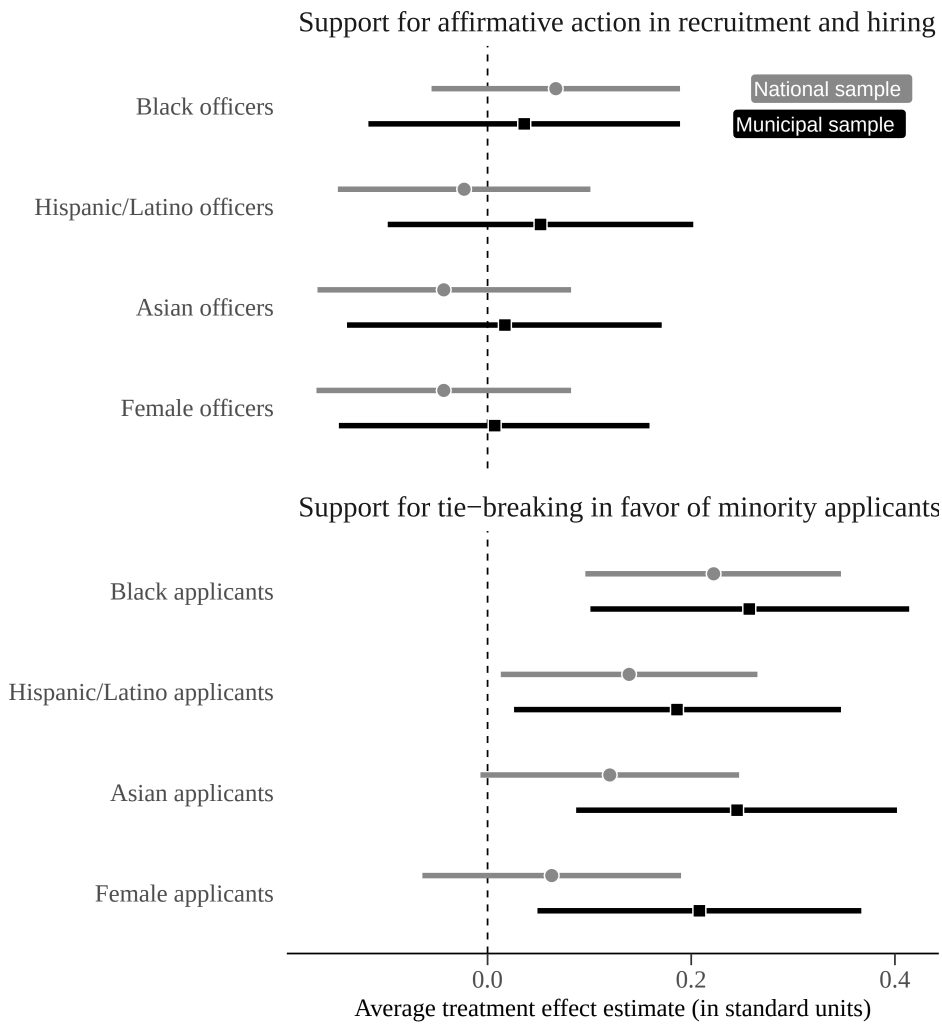

Figure S11: Estimated treatment effects on survey index components in each sample. Point estimates for the ATEs estimated using OLS regression of the outcome on treatment assignment, with robust standard errors in parentheses. All estimates are standardized using Glass's Δ, which scales outcomes by the standard deviation in the control group [28, 29].

S2.1.2 Average treatment effects estimated with regression adjustment

Our pre-analysis plan (PAP) for the experiment on the municipal sample specified that we would estimate Average Treatment Effects (ATEs) using regression adjustment with pre-treatment covariates to increase precision (see Section S3). However, the PAP was not explicit regarding the specific subset of covariates that would be used for regression adjustment across outcome measures, and different subsets may be more or less prognostic of different outcome measures. Given this ambiguity, we opted for a conservative decision to report estimates without regression adjustment in the manuscript unless there was a pre-treatment measure of that outcome available from the baseline survey (conducted approximately 4 months prior to the experiment).

For two of the outcomes in the municipal experiment (trust and confidence in the police, and willingness to cooperate) pre-treatment measures of the exact same outcomes were available from the baseline survey. The estimates reported in the manuscript for these two outcomes are from the standard covariate-adjusted linear regression estimator [30] that interacts treatment assignment with pre-treatment covariates (here, simply the baseline measure of the outcome). For all other outcomes, we reported the simple difference in means between treatment and control, estimated using Ordinary Least Squares (OLS) regression of the outcome on treatment assignment.

Here, we compare the estimates from the manuscript to those obtained via regression adjustment. To do so, we take an agnostic approach that automates the process using adaptive specification search via machine learning and cross-validation. Specifically, we apply a random-forest based cross-estimation procedure for regression adjustment using the default "out-of-the-box" settings from the `crossEstimation` package in `R` [31]. This approach yields finite-sample-unbiased estimates of the sample average treatment effects via non-parametric regression adjustments with the random forests algorithm. A key advantage of this approach is that it minimizes the risks of researcher discretion that can arise from specification search (i.e., picking different subsets of covariates for regression adjustment). This approach is ideal for our setting since it automates the process of covariate selection and estimation in regression adjustment without compromising statistical inference.

We identified a broad set of 33 pre-treatment covariates as candidate variables for regression-adjustment in the municipal survey and 17 for the national survey (which did not include a baseline wave). These lists include the demographic measures described in Tables S1-S2, as well as additional pre-treatment attitudinal measures (e.g., baseline measures of trust and cooperation in the municipal survey; pre-treatment measures of trust/confidence in the national survey). Table S5 reports the underlying point estimates and standard errors for the results presented in the manuscript (OLS estimator) alongside those obtained from regression adjustment via the random-forest based cross-estimation procedure (random forest estimator). We find limited precision gains from regression-adjustment in this application. We provide a description of these additional variables below (see also Tables S6-.S7).

Outcome measure	**Estimator**	
Sample	OLS	Random forest
Support for affirmative action in recruitment and hiring:		
Municipal sample	0.03 (0.08)	-0.01 (0.06)
National sample	0.00 (0.06)	-0.02 (0.06)
Support for tie-breaking in favor of minority applicants:		
Municipal sample	0.26 (0.08)*	0.24 (0.07)*
National sample	0.17 (0.06)*	0.13 (0.06)*
Voted for police department diversification:		
Municipal sample	0.22 (0.09)*	0.22 (0.09)*
Donation to Black police officer association:		
Municipal sample	-0.03 (0.08)	-0.04 (0.07)
Trust and confidence in the police department:		
Municipal sample	-0.14 (0.05)*	-0.14 (0.05)*
Willingness to cooperate with police officers:		
Municipal sample	0.03 (0.06)	0.03 (0.06)

Table S5: Estimated treatment effects with and without covariate adjustment. The first column of results shows point estimates for the average treatment effect (ATE) from the Ordinary Least Squares (OLS) regression estimator, with robust standard errors in parentheses. The next column of results shows covariate-adjusted point estimates (standard errors) for the same ATEs, estimated using the random-forest based cross-estimation procedure [31]. All estimates are standardized using Glass's Δ, which scales outcomes by the standard deviation in the control group [28, 29]. Only the first two outcomes were measured in the national survey. $^{*}P < 0.05$

Description of additional pre-treatment measures:

- **Number of police officers known**: "How many police officers do you know, at least as acquaintances?" [1 = None, 2 = One, 3 = Two, 4 = Three, 5 = Four, 6 = Between 5 and 9, 7 = 10 or more].
- **Frequency of contact with police**: "How often do you interact with Yonkers police?" [1 = "Never", 2 = "Less than once a year", 3 = "Yearly", 4 = "A few times a year", 5 = "Monthly", 6 = "Weekly", 7 = "Daily"].
- **Any contact with police in last 12 mos:** "During the past 12 months, have you had any contact with an officer from YPD?" [1 = "Yes", 0 = "No"].
- **Any prior arrest by police**: "Have you ever been arrested by the Yonkers Police?" [1 = "Yes", 0 = "No"].
- **Any prior unfair treatment by police**: "Have you ever been treated unfairly by the Yonkers Police?" [1 = "Yes", 0 = "No"].
- **Feelings of safety in local area**: "Generally speaking, how safe do you feel walking alone at night within a mile of where you live?" [1 = "Not at all safe", 2 = "Slightly safe", 3 = "Moderately safe", 4 = "Very safe", 5 = "Extremely safe"]

- **Any prior crime victimization**: "While living in Yonkers, have you ever been the victim of a crime?" [1 = Yes, 0 = No].
- **Partisanship**: measured using the 7-point branching question from the American National Election Studies (ANES) Survey. [1 = "Strong Democrat", 2 = "Not very strong Democrat", 3 = "Lean Democrat", 4 = "Independent", 5= "Lean Republican", 6= "Not very strong Republican", 7 = "Strong Republican"]
- **Attentiveness to local/national politics**: "How often do you pay attention to what's going on in government and politics at the [local/national] level?" [1 = "Never", 2 = "Some of the time", 3 = "About half the time", 4 = "Most of the time", 5 = "Always"]
- **Trust and confidence in police (2-item index)**: The 2-item trust and confidence measure described in Section S1.3.1 was also measured in the baseline survey.
- **Willingness to cooperate with police (4-item index)**: The 4-item cooperation measure described in Section S1.3.1 was also measured in the baseline survey.
- **Legitimacy, trust, and confidence (10-item index)**: The two items from the trust and confidence measure as well as the eight items listed below were combined to create a 10 item index ($\alpha = 0.95$). Responses to each item below were recorded using a 7 point scale from "Strongly disagree" to "Strongly agree" with a neutral midpoint. Prompt: "Please say whether you agree or disagree with the below statements about the police in Yonkers." Items were presented in random order.
    1. "They care about the well-being of people they deal with"
    2. "They make fair and impartial decisions"
    3. "They stand up for values that are important to you"
    4. "They behave according to the law when dealing with people"
    5. "They make me feel safer in my neighborhood"
    6. "They treat people equally"
    7. "They are trying to make my community better"
    8. "They respect the people in my community"
- **Stated support for diversification policy**. Respondents were first asked their support for four potential police policy changes (order randomized): diversification, civilian oversight, body worn cameras, and community policing. Prompt: "There are ongoing discussions at the national level about a variety of policy changes that seek to improve police-community relations in one way or another. Please consider the policies described below, and whether you support or oppose them being implemented at the Yonkers Police Department." Each potential policy change (as described below) was displayed in random order, and responses were recorded using a 7-point scale from "Strongly oppose" (1) to "Strongly support" (7) with a neutral midpoint.
    1. **Diversification**. Implement affirmative actions programs that increase recruitment and hiring of officers from underrepresented groups so that police more closely resemble the community in terms of race/ethnicity and gender.
    2. **Civilian oversight**. Create a Civilian Review Board with the power to investigate and recommend action for complaints made against police officers for instances that include excessive force, abuse of authority, and offensive language.

3. **Body worn cameras**. Require police officers to wear body cameras that record their interactions with the public while on duty.
4. **Community policing**. Establish regular meetings between police and the public that provide a forum for city representatives, businesses, and residents to share information and cooperatively address neighborhood issues.

- **Rank ordering of diversification policy**. Respondents were asked to rank order the relative importance of each of the four policy changes listed above being implemented in their local police department (see Fig. S12). Responses were re-coded so that 1 indicates the least preferred policy change and 4 indicates the most preferred policy change.

- **Trust and confidence in police (6-item index)**: A six-item trust and confidence index ($\alpha = 0.91$) developed by Pew Research [32]. These questions were only asked in the national sample. Respondents were asked each of the six questions below in randomized order, with responses recorded on a 5 point scale: 1= "Never", 2 = "Rarely", 3 = "Sometimes", 4 = "Often", 5 = "Always". Prompt: "In your view, how much of the time do police officers ..."

1. Care about people like you
2. Do a good job protecting people from crime
3. Handle the resources available to them in a responsible way
4. Provide fair and accurate information to the public
5. Admit mistakes and take responsibility for them
6. Treat racial and ethnic groups equally

Please consider each of the policies described below and rank the relative importance, in your view, of them being implemented at the Yonkers Police Department.

The "most important" should be at the top of the list ("1") and the "least important" should be at the bottom of the list ("4").

Body worn cameras. Require police officers to wear body cameras that record their interactions with the public while on duty.

Civilian oversight. Create a Civilian Review Board with the power to investigate and recommend action for complaints made against police officers for instances that include excessive force, abuse of authority, and offensive language.

Community policing. Establish regular meetings between police and the public that provide a forum for city representatives, businesses, and residents to share information and cooperatively address neighborhood issues.

Diversification. Implement affirmative actions programs that increase recruitment and hiring of officers from underrepresented groups so that police more closely resemble the community in terms of race/ethnicity and gender.

Figure S12: Rank ordering of preferences for police policy change question

	Mean	SD	Min	Max	N
Predicted - actual share of US police officers by race/ethnicity and sex:					
White	-13.56	21.64	-72.00	28.00	998
Black	7.04	13.78	-11.00	89.00	998
Hispanic/Latino	1.27	11.85	-13.00	87.00	998
Asian or another race/ethnicity	5.26	9.28	-4.00	96.00	998
Male	-21.49	17.00	-88.00	12.00	998
Female	21.49	17.00	-12.00	88.00	998
Additional pre-treatment measures from national survey:					
Frequency of contact with police	2.42	1.50	1.00	7.00	997
Any contact with police in last 12 mos	0.37	0.48	0.00	1.00	998
Feelings of safety in local area	3.18	1.28	1.00	5.00	998
Partisanship	3.58	2.31	1.00	7.00	997
Trust and confidence in police (6-item index)	0.00	0.82	-2.12	1.59	998

Table S6: Additional pre-treatment measures used for regression adjustment in the national sample. Here we present descriptive statistics for the subset of respondents assigned to receive either the information treatment or control (N = 998) rather than the full sample (N= 2,017), which includes the two additional treatment arms not assigned in the municipal sample.

	Mean	SD	Min	Max	N
Prior experience with police and crime in baseline survey:					
Number of police officers known	2.54	1.91	1.00	7.00	644
Frequency of contact with police	2.38	1.31	1.00	7.00	644
Any contact with police in last 12 mos	0.41	0.49	0.00	1.00	644
Any prior arrest by police	0.04	0.19	0.00	1.00	643
Any prior unfair treatment by police	0.13	0.34	0.00	1.00	644
Feelings of safety in local area	3.16	1.10	1.00	5.00	644
Any prior crime victimization	0.34	0.47	0.00	1.00	644
Additional background measures from baseline survey:					
Currently employed	0.60	0.49	0.00	1.00	644
Not currently employed	0.21	0.41	0.00	1.00	644
Retired	0.19	0.39	0.00	1.00	644
Duration of residency (years)	23.43	17.20	0.00	82.00	644
Homeowner	0.62	0.49	0.00	1.00	644
Partisanship	2.90	1.94	1.00	7.00	644
Attentiveness to local politics	3.28	1.13	1.00	5.00	643
Attentiveness to national politics	4.00	0.95	1.00	5.00	643
Trust and confidence in police (2-item index)	0.00	0.93	-2.45	1.60	644
Willingness to cooperate with police (4-item index)	0.00	0.76	-2.96	0.87	644
Legitimacy, trust, and confidence (10-item index)	0.00	0.86	-2.26	1.54	644
Stated support for diversification policy	5.51	1.75	1.00	7.00	644
Rank ordering of diversification policy	2.25	1.04	1.00	4.00	644
Predicted - actual share of local police officers by race/ethnicity and sex:					
White	-18.46	17.73	-78.00	22.00	644
Black	9.49	8.14	-6.00	44.00	644
Hispanic/Latino	4.95	11.34	-15.00	85.00	644
Asian	4.01	4.37	-1.00	30.00	644
Female	11.99	13.63	-13.00	78.00	644
Male	-11.99	13.63	-78.00	13.00	644
Additional pre-treatment measures from followup survey:					
Any police contact since baseline survey	0.19	0.39	0.00	1.00	644
Any unfair treatment by police since baseline	0.03	0.18	0.00	1.00	642
Feelings of safety in local area	3.07	1.04	1.00	5.00	644
Victim of crime since baseline survey	0.03	0.18	0.00	1.00	644

Table S7: Additional pre-treatment measures used for regression adjustment in the municipal sample. For the measures captured in the baseline survey (N = 1,413), we only present descriptive statistics for the subset of individuals that also completed the followup survey (N = 644). The predicted - actual shares of local police officers by race/ethnicity and sex were measured in the followup survey after the other pre-treatment measures listed in the table, and prior to treatment assignment.

S2.1.3 Conditional average treatment effects estimated with causal forests

In this section, we examine treatment effect heterogeneity in the information provision experiment as a function of respondents' pre-treatment covariates. Our PAP for the experiment on the municipal sample (see S3) specified that we would estimate Conditional Average Treatment Effects (CATEs) for sub-groups of respondents defined by race/ethnicity, sex, partisanship, and pre-treatment measures of belief accuracy about the race/ethnicity and gender composition of police officers. Additionally, we specified that we would conduct a broader exploratory search for treatment effect heterogeneity as a function of pre-treatment covariates using causal forests. Here, we automate the search for treatment effect heterogeneity using causal forests, an implementation of the Generalized Random Forests (GRF) algorithm which estimates heterogeneity as a function of respondents' background covariates, and generates individual-level predictions of causal effects for the entire sample [33, 34, 35]. We implement this via the `grf` package for `R`, using the recommended default settings with honest splitting and 4000 trees [34, 36]. For these analyses, we use the same set of pre-treatment covariates that were used for regression adjustment, as described in Section S2.1.2.

Following graphical presentations in prior work [37, 38], Figures S13-S14 plot the causal forest estimated treatment effects (and 95% CIs) for each individual as a function of their covariate profiles to provide an overall summary of treatment effect heteogeneity across outcome measures. These visual summaries show little evidence of treatment effect heterogeneity. Table S8 provides results from the omnibus test of treatment effect heterogeneity using the "best linear predictor" method proposed by Athey and Wager (2019), Section 2.2 [34]. Briefly (see [34, 36] for details), this procedure tests whether heterogeneity in the out-of-bag causal forest estimates, denoted $\hat{\tau}^{(-1)}(X_i)$, is associated with heterogeneity in the CATE function, $\tau(X_i)$.[2] This test is performed via OLS regression of a transformed outcome that represents predicted treatment effects in the held out dataset, denoted $\tilde{Y}_i$, on C_i and D_i, as defined below:

- $\tilde{Y}_i = Y_i - \hat{m}^{(-i)}(X_i)$. Y_i denotes the observed outcome vector and $\hat{m}^{(-i)}(X_i)$ denotes the vector of out-of-bag estimates for the expected outcome, marginalizing over treatment (i.e., $m(x) = \mathbb{E}[Y_i|X_i = x]$ for binary treatment Z_i and covariates X_i).
- $C_i = \bar{\tau}\left(Z_i - \hat{e}^{(-i)(X_i)}\right)$. $\bar{\tau}$ denotes the average of the out-of-bag treatment effect estimates and $\hat{e}^{(-i)}(X_i)$ denotes the out-of-bag estimates for the propensity score (i.e., $e(x) = \Pr(Z_i|X_i = x)$).
- $D_i = \left(\hat{\tau}^{(-i)}(X_i) - \bar{\tau}\right)\left(Z_i - \hat{e}^{(-i)}(X_i)\right)$, where $\hat{\tau}^{(-i)}(X_i)$ again denotes the out-of-bag casual forest estimates for each individual.

The coefficient on D_i is then interpreted as a measure of the quality of the causal forest estimates of treatment effect heterogeneity [34]. If the coefficient on D_i is 1 then the estimates of treatment effect heterogeneity are "well calibrated," but if significant and positive this provides evidence of an association between $\hat{\tau}^{(-1)}(X_i)$ and $\tau(X_i)$. The estimated coefficients for D_i (with robust standard errors in parenthesis) are provided for each outcome in Table S8 alongside t-statistics and one-sided P-values for omnibus test – if the estimated coefficient is significantly

[2] If the CATE function is constant then $\tau(x) = \mathbb{E}[Y_i(1) - Y_i(0)|X_i = x] = \mathbb{E}[Y_i(1) - Y_i(0)] = \tau$.

greater than 0 then we reject the null of treatment effect homogeneity. We fail to reject the null of treatment effect homogeneity for all outcomes, across both samples.

Table S9 provides another overall summary of the individual level predictions plotted in Figures S13-S14, grouped by outcome and sample. For those outcomes in which the estimated ATE was statistically distinguishable from zero, we also find that individual-level predictions are in the expected direction. First, support for tie-breaking in favor of minority applicants: 100% of the estimates were positively signed in the municipal sample (99% in national sample); 64% of the 95% CIs excluded zero in the municipal sample (25% in national sample); among those estimates with CIs that excluded zero, all were positively signed in both samples. Second, voting for police department diversification (municipal sample only): 100% of the estimates were positively signed; 33% of the 95% CIs excluded zero; among those estimates with CIs that excluded zero, all were positively signed. Third, trust and confidence in the police (municipal sample only): 100% of the estimates were negatively signed; 58% of the 95% CIs excluded zero; among those estimates with CIs that excluded zero, all were negatively signed.

Finally, we plot the causal forest estimates for respondents' CATEs, $\hat{\tau}^{(-i)}(X_i)$, against the subset of covariates that we pre-registered an intention to provide estimated CATEs for in Figures S15-S32. Figures S15-S26 plot respondents' $\hat{\tau}^{(-i)}(X_i)$ estimates against each measure of belief accuracy, defined as the difference between their pre-treatment guess about the share of officers in demographic sub-group and the actual share. For example, Figure S17 shows the relationship between respondents' causal forest estimated treatment effects (vertical axis) and their belief accuracy (predicted - actual share of Black police officers) in the municipal sample. Here we see some evidence that treatment effects were moderated by belief accuracy. This suggests, for example, that effects on support for tie-breaking in favor of minority applicants were somewhat stronger among those that overestimated the share of Black officers at YPD.

We caution, however, that the overall picture is unclear. Although the point estimates are consistent with some moderation by belief accuracy we cannot reject the null hypothesis of treatment effect homogeneity, possibly due to sample size constraints. There is much weaker evidence for heterogeneity across sub-groups defined by race/ethnicity (Fig S27-S28), sex (Fig. S29-S30), or partisanship (Fig. S31-S31). In short, the causal forest estimates do not reliably identify sub-groups of respondents for which evidence of stronger (or weaker) treatment effects is obvious. Any treatment effect heteogeneity that may be present seems relatively weak, and limited to the pre-treatment measures of belief accuracy. Ultimately, sample size constraints limit our ability to reliably detect small but potentially meaningful variation in causal effects as a function of respondents' belief accuracy.

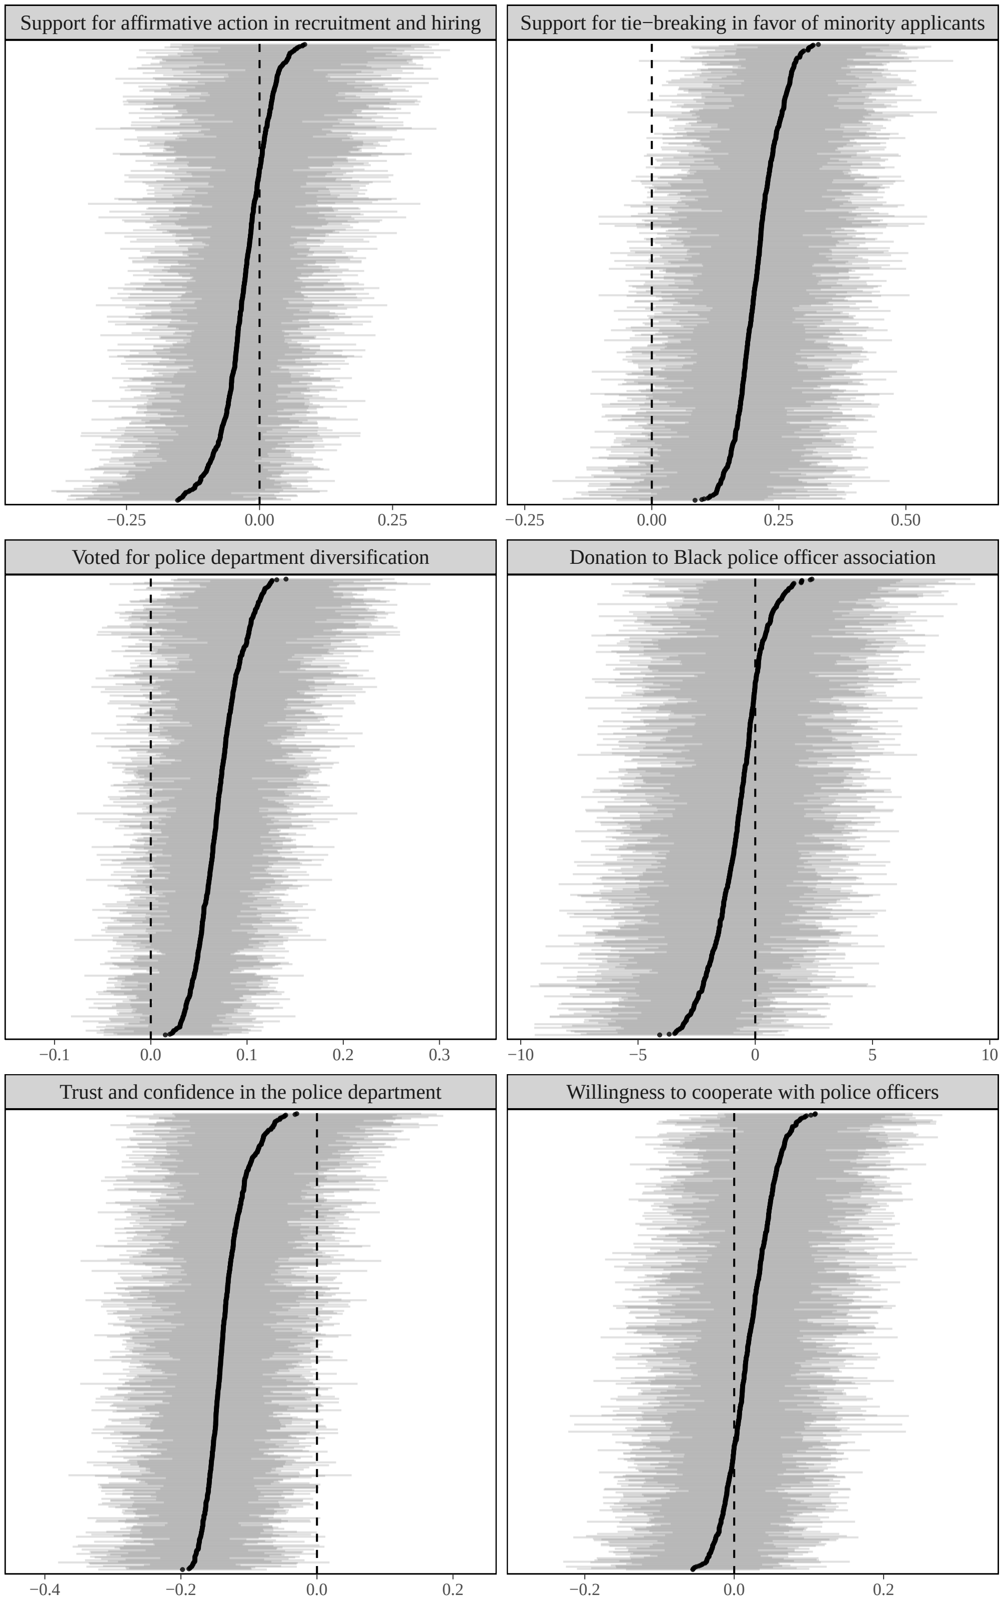

Figure S13: Causal forest estimated treatment effects in municipal sample by outcome measure. Estimated treatment effects for each individual as a function of their covariate profile (black dots) and 95% confidence intervals (grey bars).

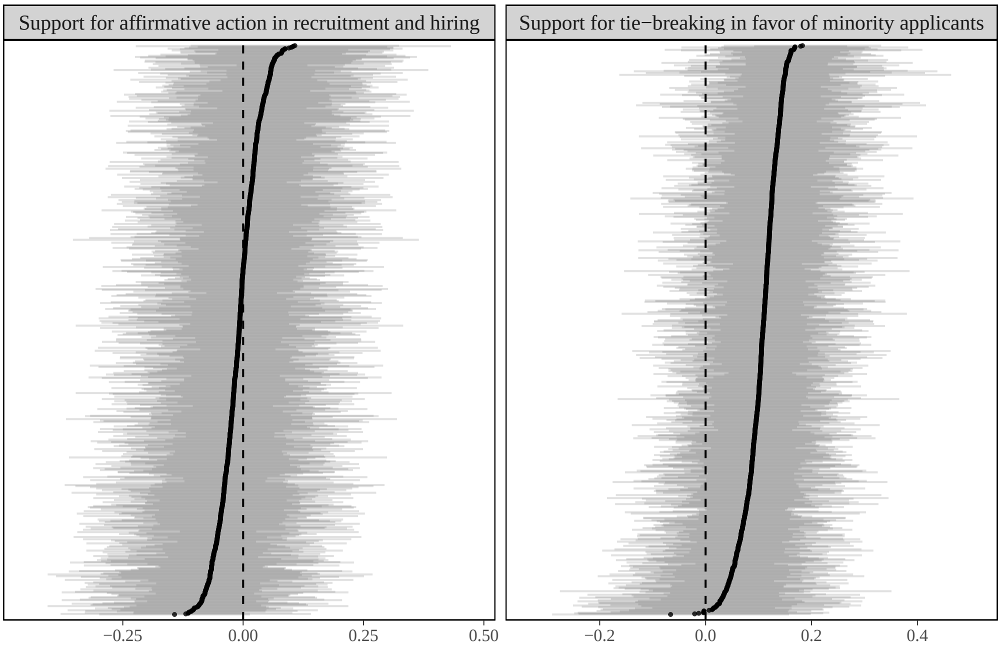

Figure S14: Causal forest estimated treatment effects in national sample by outcome measure. Estimated treatment effects for each individual as a function of their covariate profile (black dots) and 95% confidence intervals (grey bars).

	Estimate	***t*-statistic**	$\mathbf{Pr}(T \geq t \mid H_0)$
Support for affirmative action in recruitment and hiring:			
Municipal sample	-2.13 (1.33)	-1.61	0.95
National sample	-4.90 (1.37)	-3.57	1.00
Support for tie-breaking in favor of minority applicants:			
Municipal sample	-2.96 (1.54)	-1.92	0.97
National sample	-8.04 (1.44)	-5.58	1.00
Voted for police department diversification:			
Municipal sample	-1.16 (1.31)	-0.88	0.81
Donation to Black police officer association:			
Municipal sample	-0.85 (1.39)	-0.61	0.73
Trust and confidence in the police department:			
Municipal sample	-3.87 (1.55)	-2.49	0.99
Willingness to cooperate with police officers:			
Municipal sample	-3.52 (1.30)	-2.72	1.00

Table S8: Results from omnibus tests for heterogeneity using the causal forest estimated treatment effects. Only the first two outcomes were measured in the national survey. Point estimates (robust standard errors in parentheses) from “best linear predictor” method described in Athey and Wager (2019) Section 2.2 [34]. One sided P-values are for the null hypothesis of treatment effect homogeneity.

	Point estimates		Confidence intervals		Significant differences	
	Pos. sign	Neg. sign	Inc. zero	Excl. zero	Pos. sign	Neg. sign
Support affirmative action						
Municipal sample	0.27	0.73	0.99	0.01	0.00	1.00
National sample	0.39	0.61	1.00	0.00	-	-
Support for tie-breaking						
Municipal sample	1.00	0.00	0.36	0.64	1.00	0.00
National sample	0.99	0.01	0.75	0.25	1.00	0.00
Voted for diversification						
Municipal sample	1.00	0.00	0.67	0.33	1.00	0.00
Donation to Black officer assoc.						
Municipal sample	0.23	0.77	0.99	0.01	0.00	1.00
Trust and confidence						
Municipal sample	0.00	1.00	0.42	0.58	0.00	1.00
Willingness to cooperate						
Municipal sample	0.73	0.27	1.00	0.00	-	-

Table S9: Summary of causal forest estimated treatment effects by outcome measure and sample. Only the first two outcomes were measured in the national survey. Proportion of significant differences by sign are omitted when all 95% confidence intervals include zero.

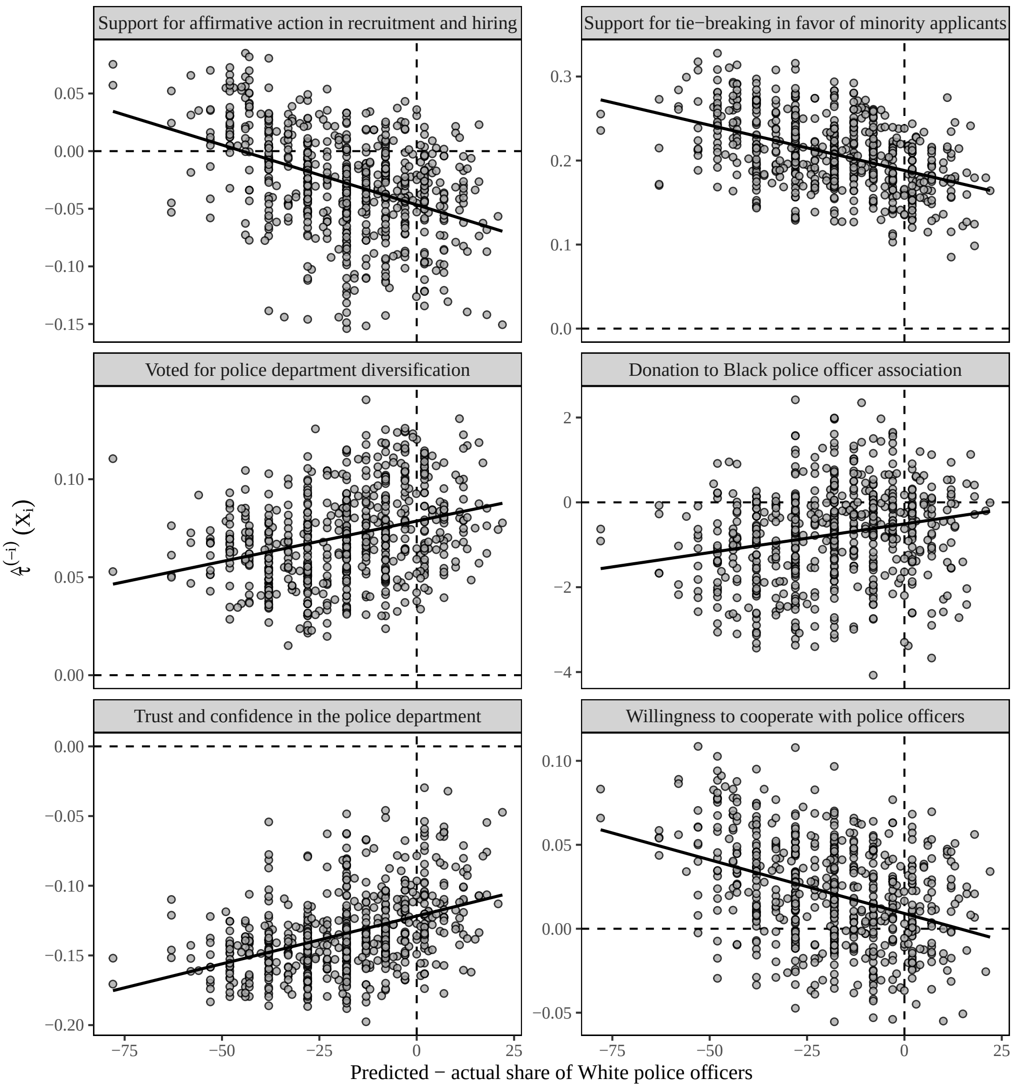

Figure S15: Causal forest estimated treatment effects by differences between predicted and actual share of White officers in municipal sample.

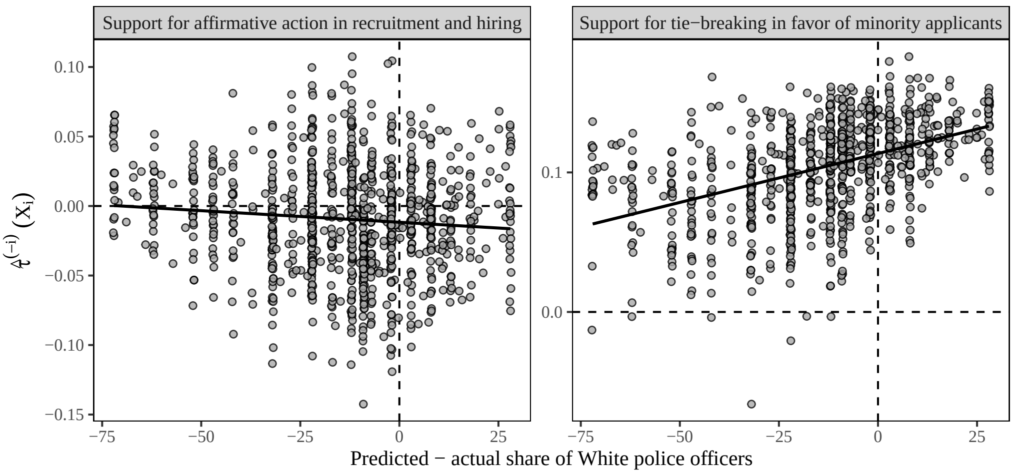

Figure S16: Causal forest estimated treatment effects by differences between predicted and actual share of White officers in national sample.

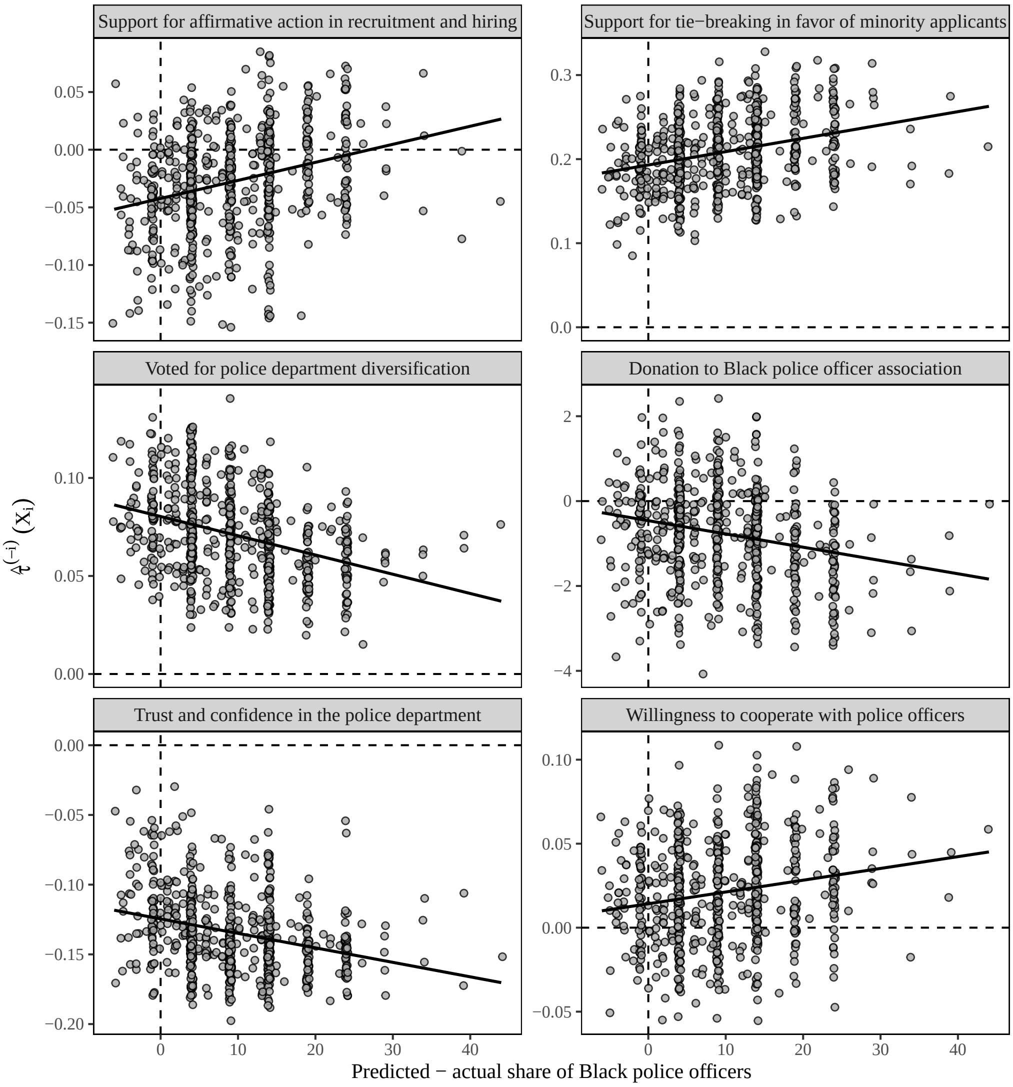

Figure S17: Causal forest estimated treatment effects by differences between predicted and actual share of Black officers in municipal sample.

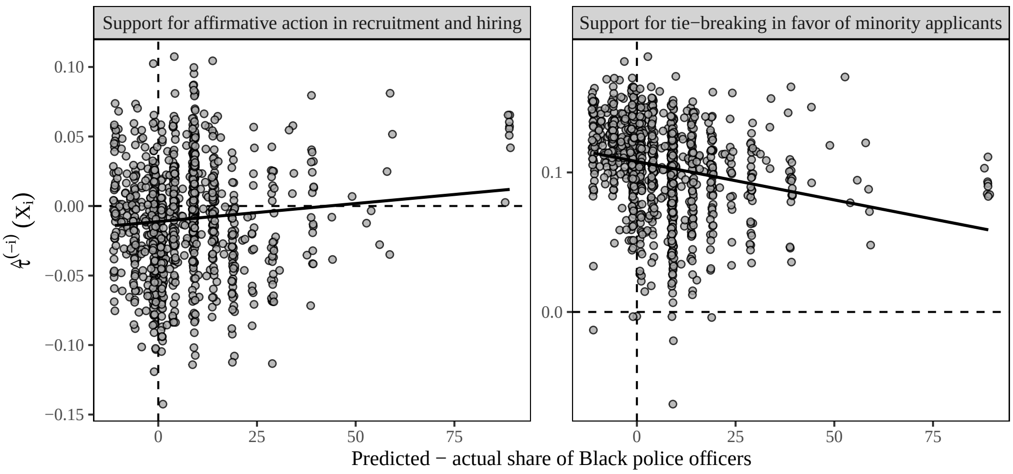

Figure S18: Causal forest estimated treatment effects by differences between predicted and actual share of Black officers in national sample.

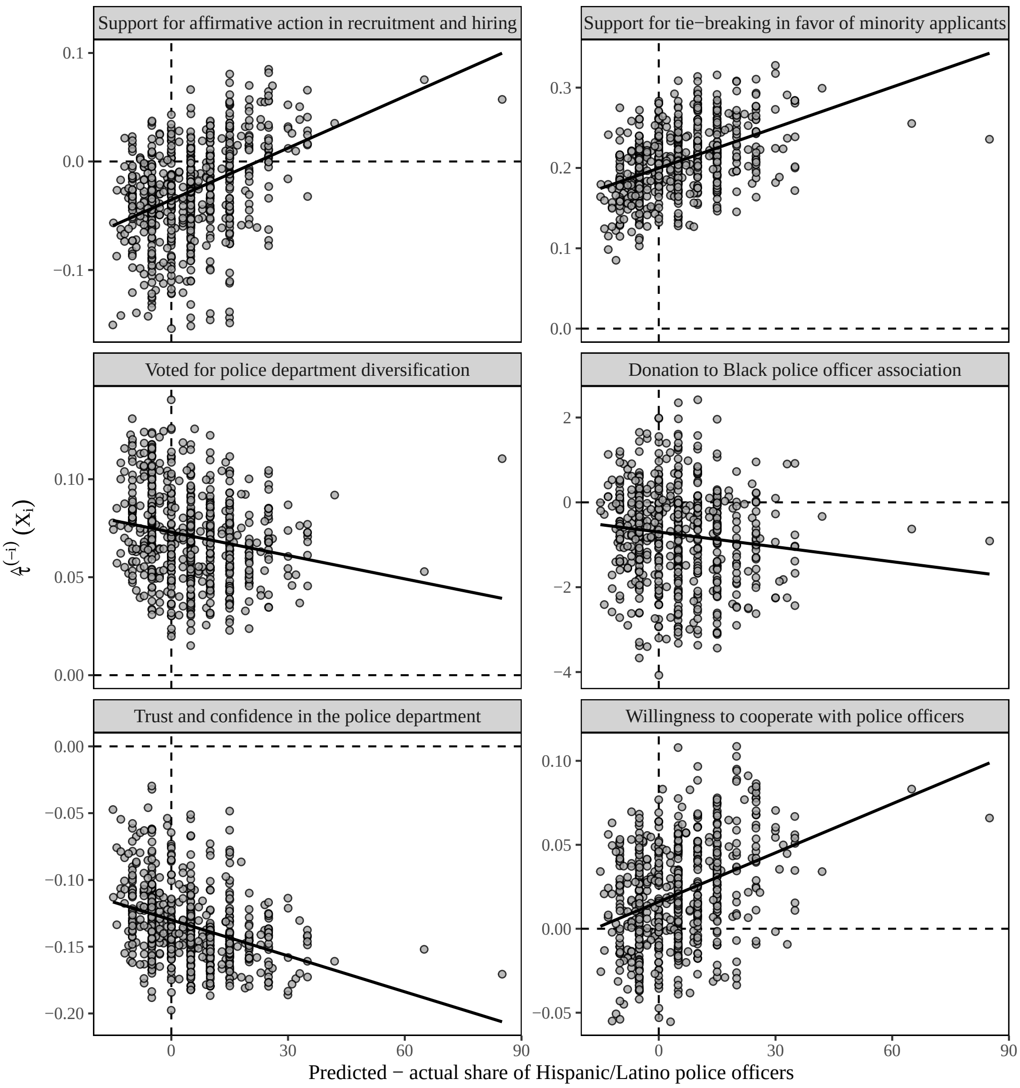

Figure S19: Causal forest estimated treatment effects by differences between predicted and actual share of Hispanic/Latino officers in municipal sample.

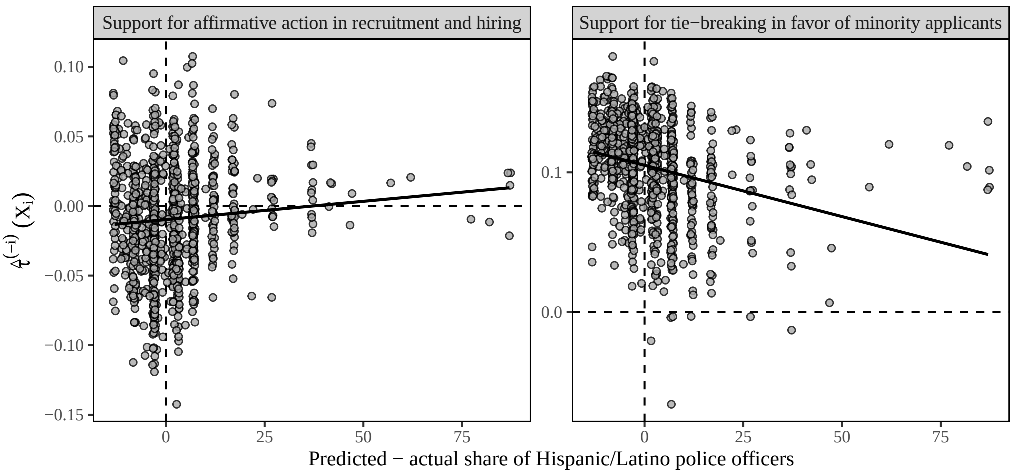

Figure S20: Causal forest estimated treatment effects by differences between predicted and actual share of Hispanic/Latino officers in national sample.

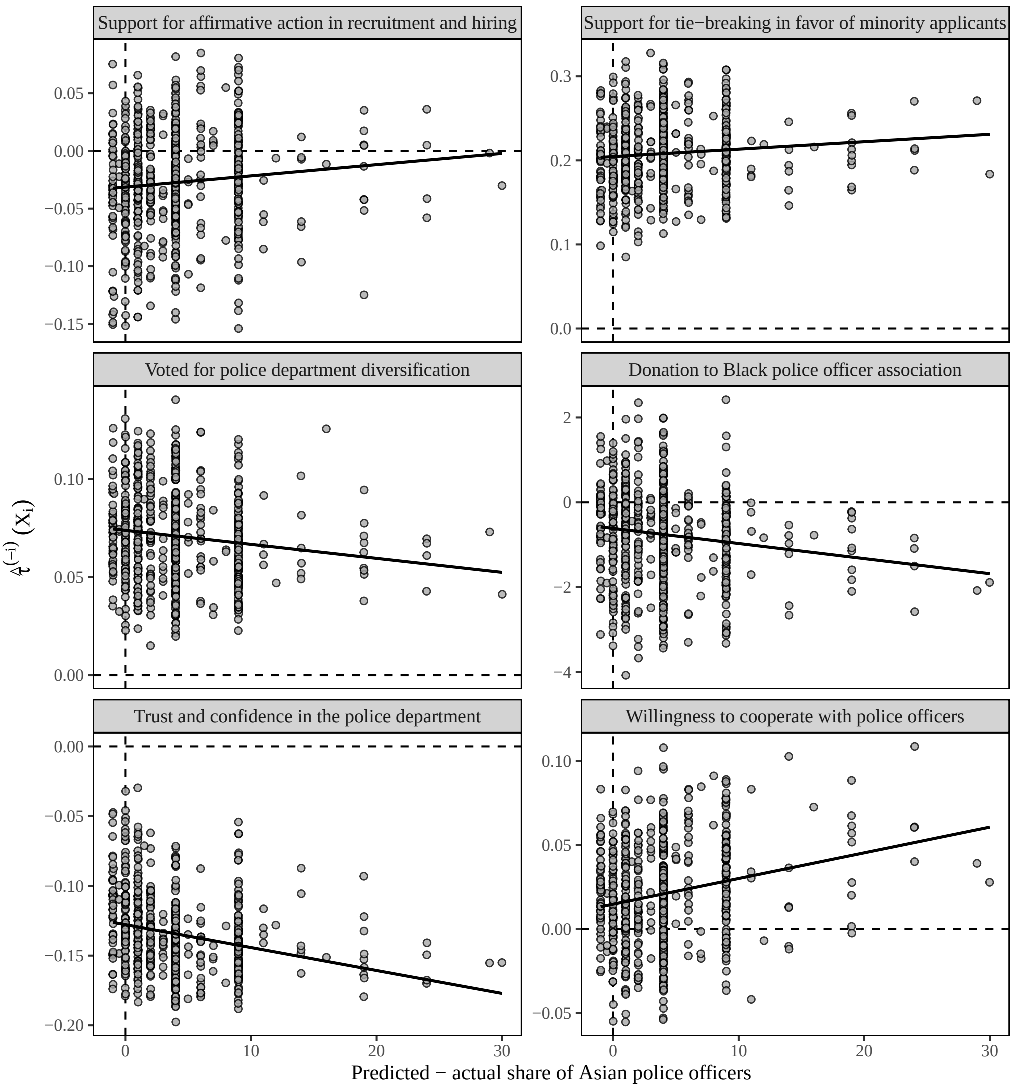

Figure S21: Causal forest estimated treatment effects by differences between predicted and actual share of Asian officers in municipal sample.

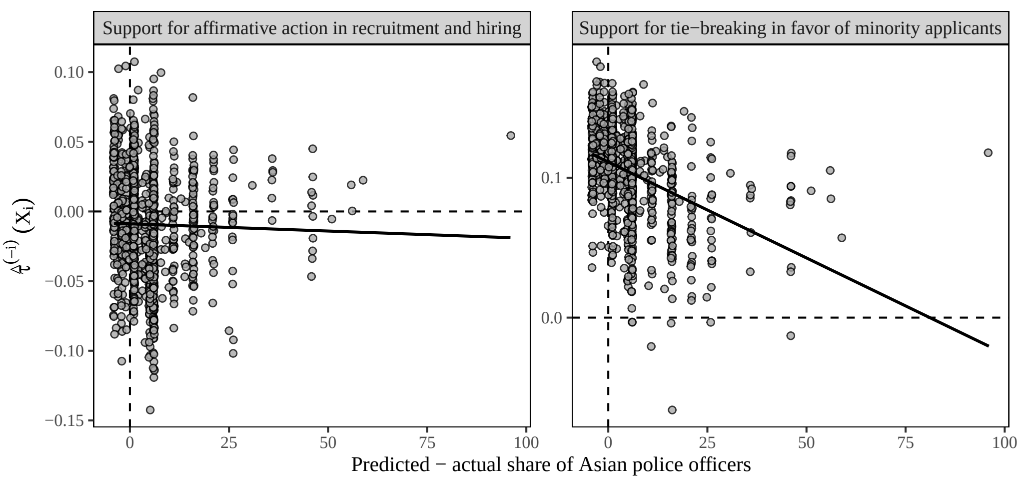

Figure S22: Causal forest estimated treatment effects by differences between predicted and actual share of Asian officers in national sample.

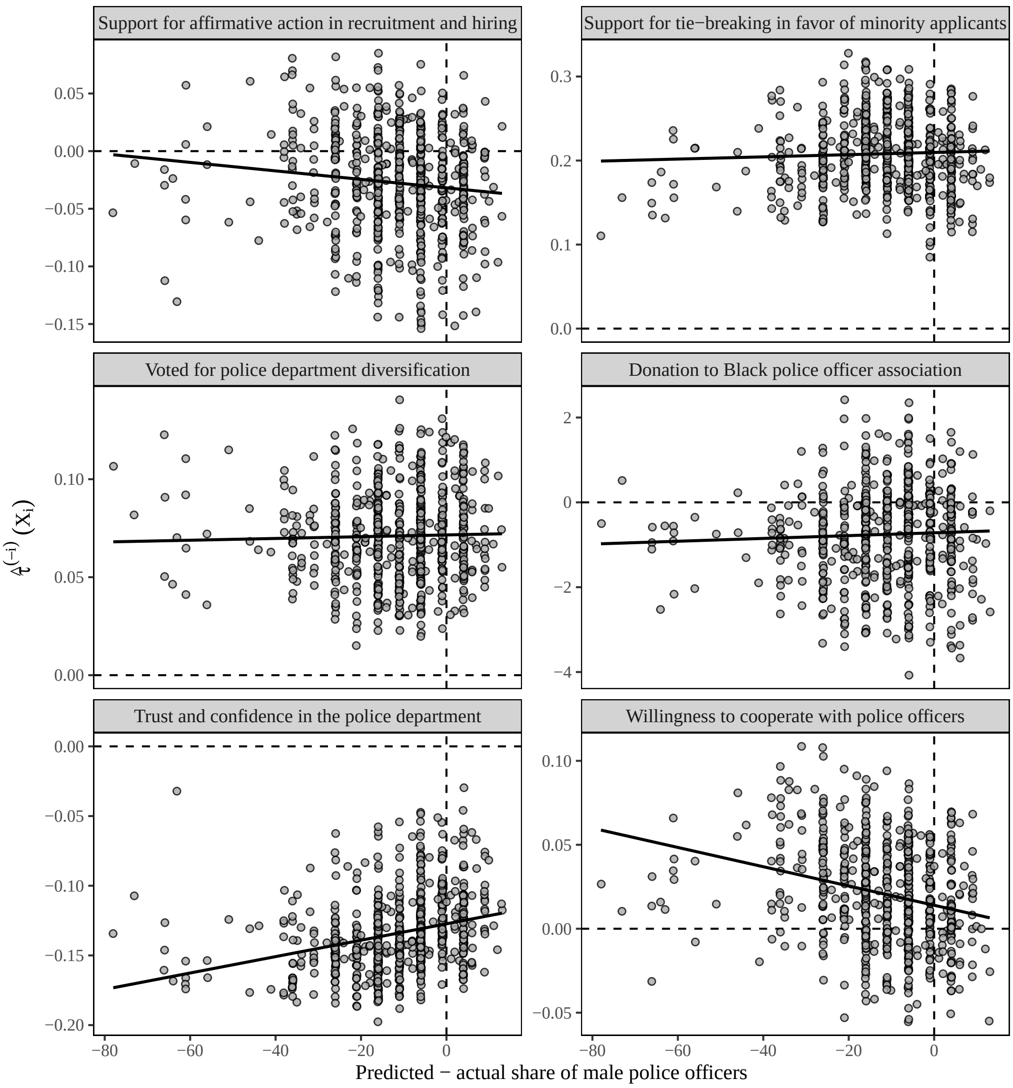

Figure S23: Causal forest estimated treatment effects by differences between predicted and actual share of male officers in municipal sample.

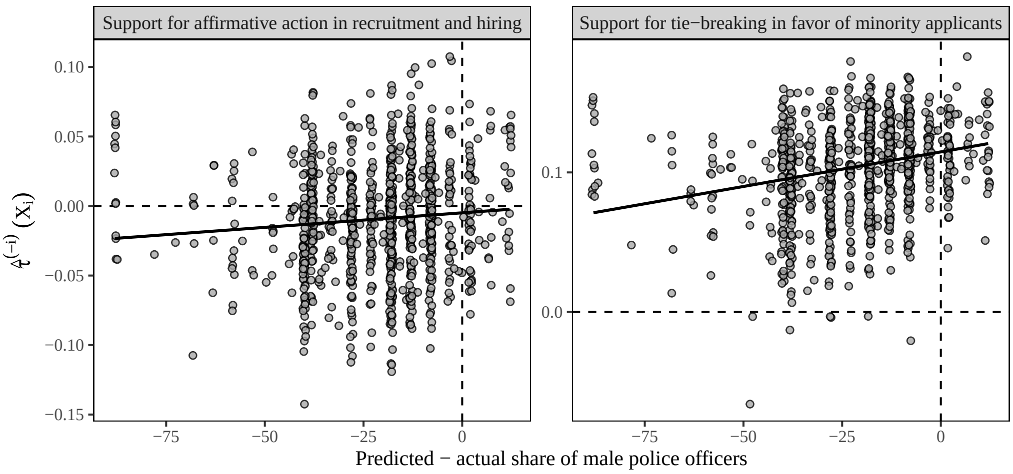

Figure S24: Causal forest estimated treatment effects by differences between predicted and actual share of male officers in national sample.

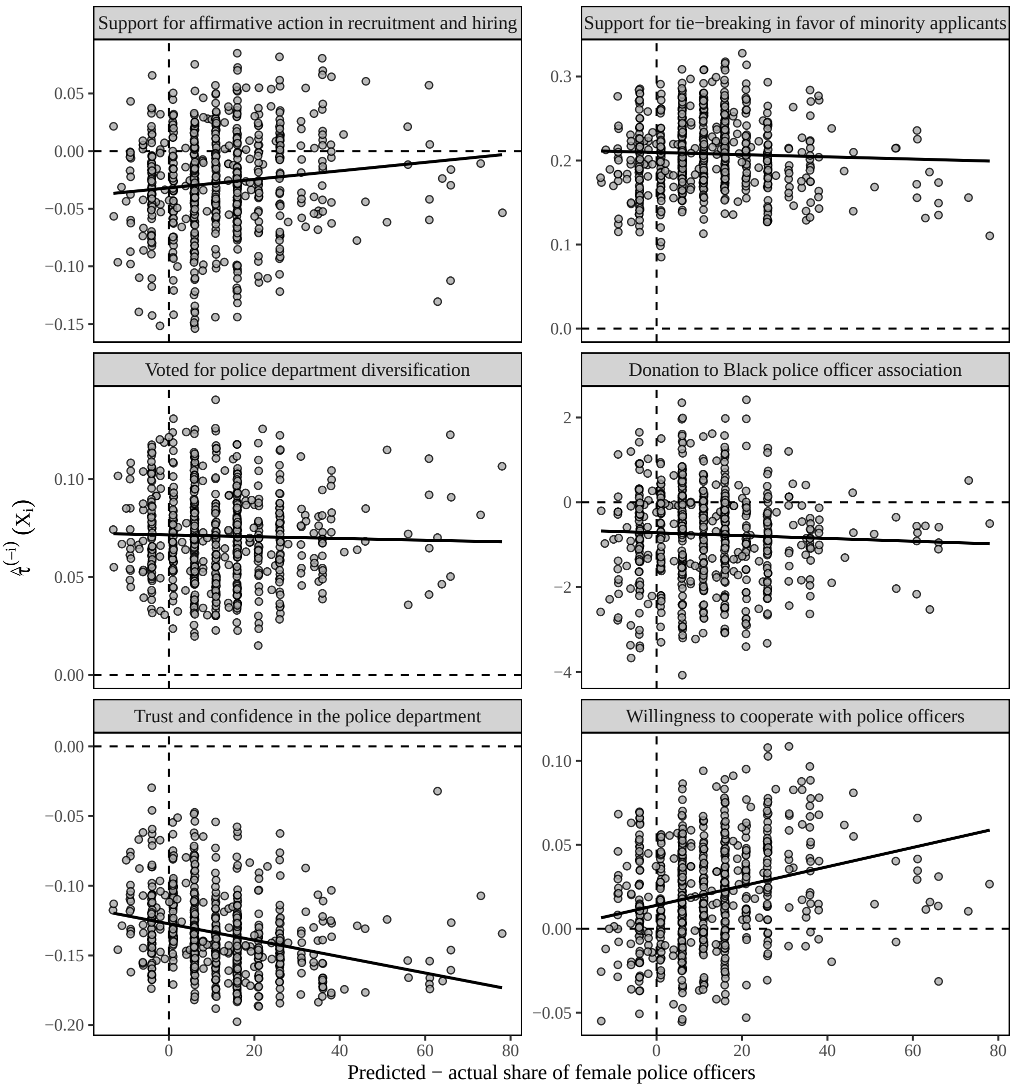

Figure S25: Causal forest estimated treatment effects by differences between predicted and actual share of female officers in municipal sample.

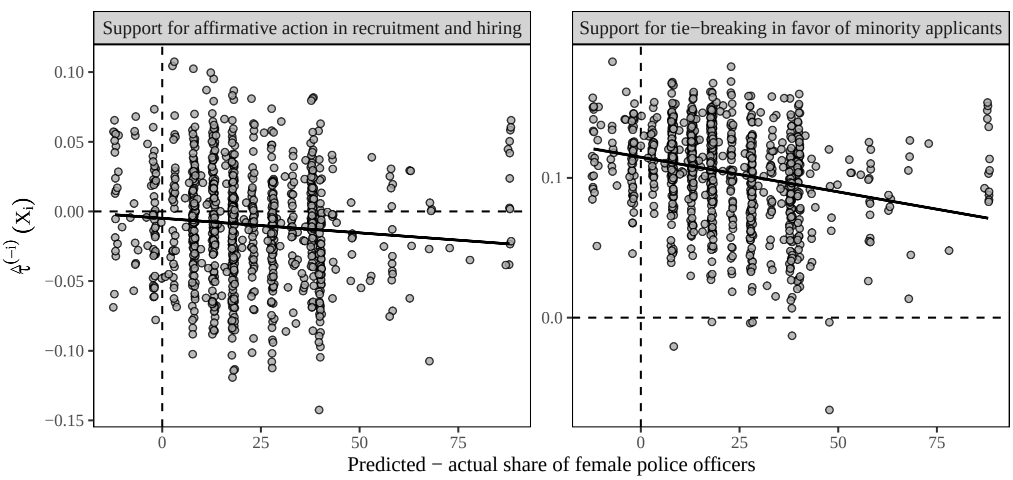

Figure S26: Causal forest estimated treatment effects by differences between predicted and actual share of female officers in national sample.

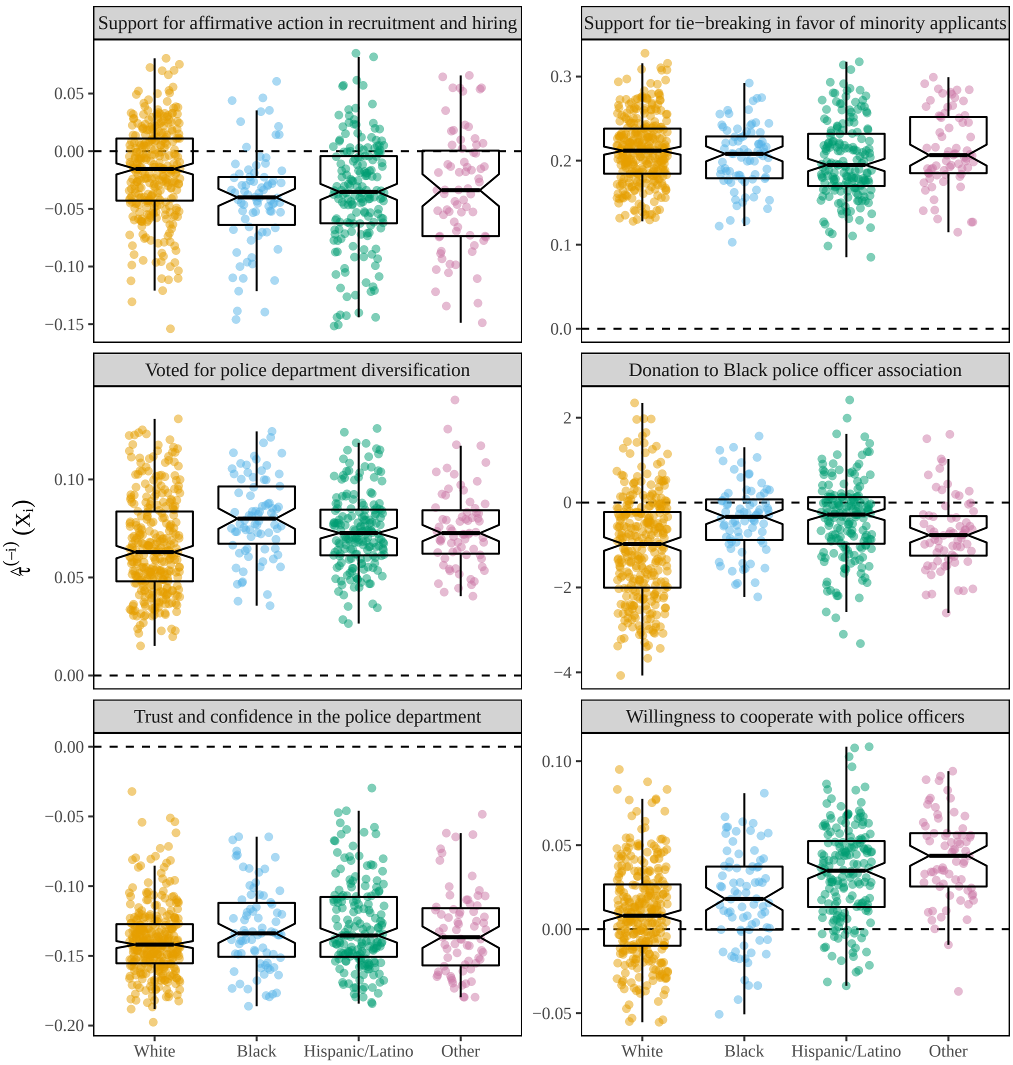

Figure S27: Causal forest estimated treatment effects by race/ethnicity in municipal sample.

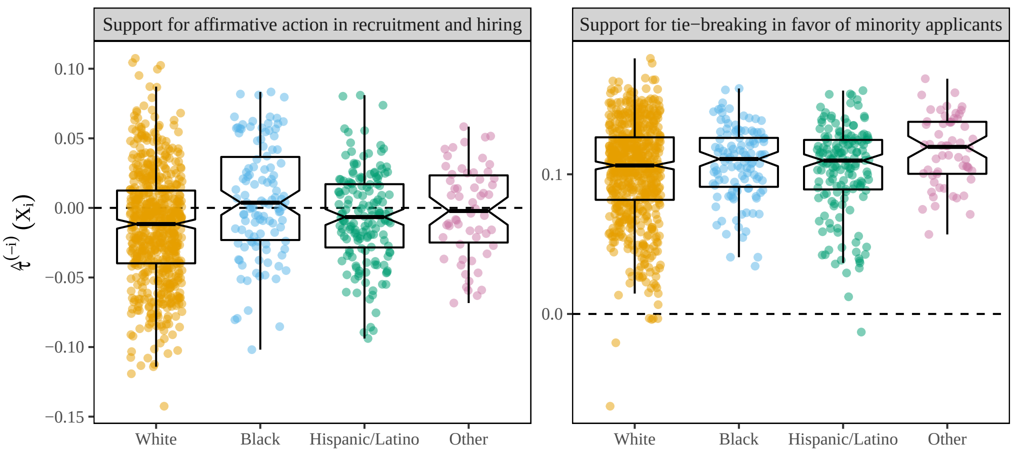

Figure S28: Causal forest estimated treatment effects by race/ethnicity in national sample.

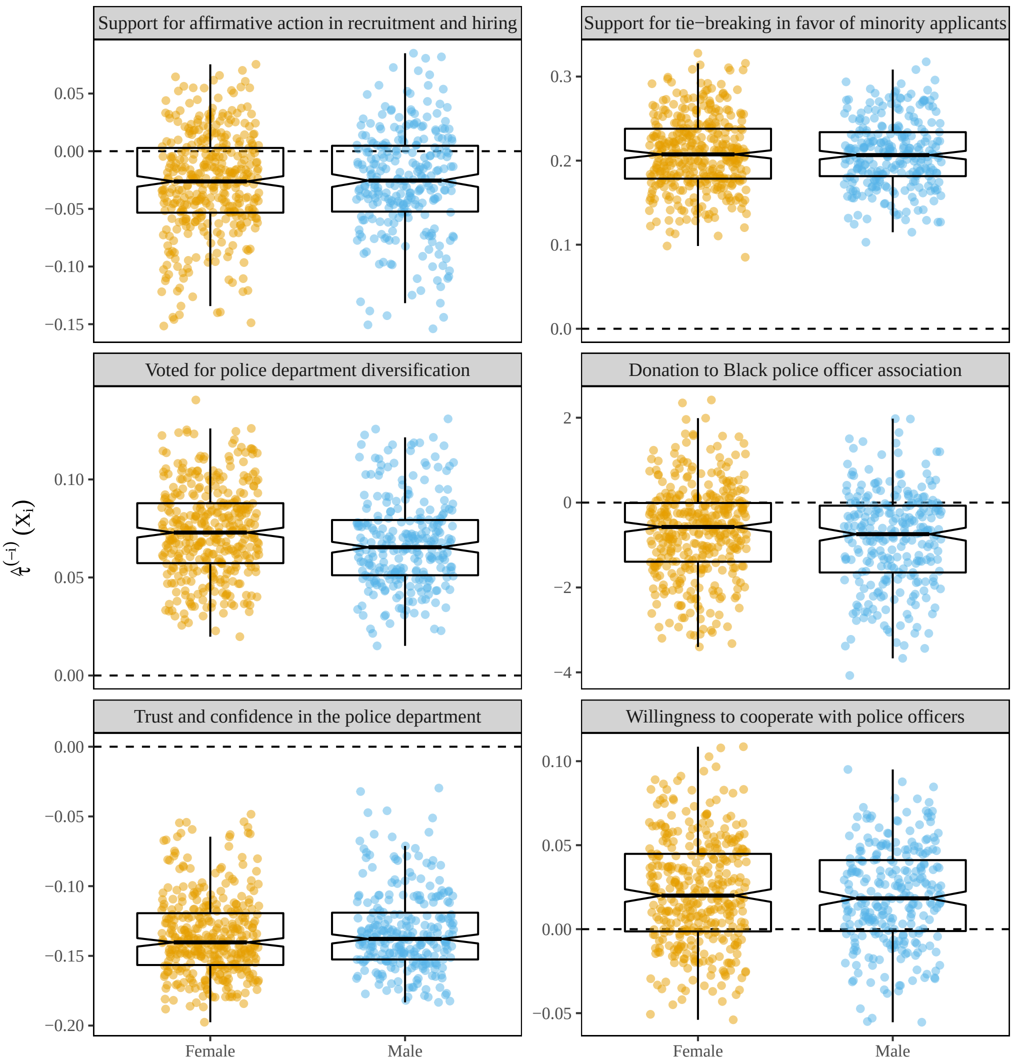

Figure S29: Causal forest estimated treatment effects by sex in municipal sample.

Figure S30: Causal forest estimated treatment effects by sex in national sample.

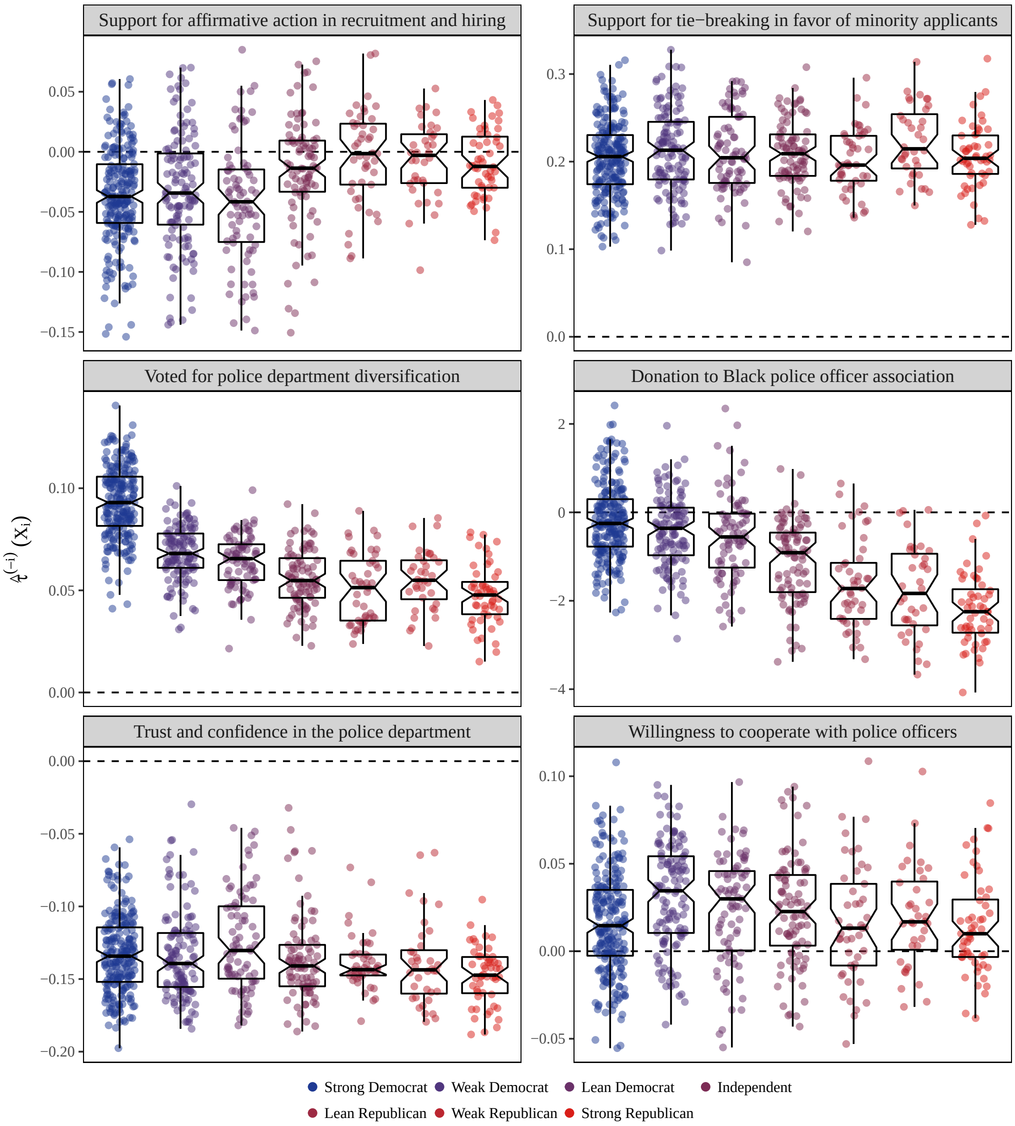

Figure S31: Causal forest estimated treatment effects by partisanship in municipal sample.

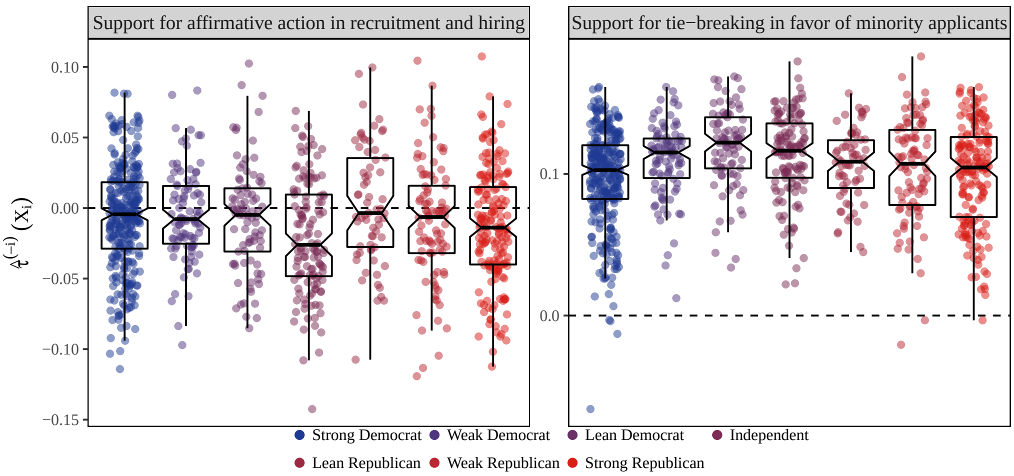

Figure S32: Causal forest estimated treatment effects by partisanship in national sample.

S2.1.4 Local causal effects estimated with instrumental variables regression

In this section, we conduct exploratory analyses to investigate the role of belief updating in explaining the effects observed in the information provision experiment. Our goal here is to obtain estimates of causal effects that are "local" to the subset of respondents who would be induced to update their beliefs if treated with information about the demographic composition of their local police department. To do so, we rely on a causal instrumental variables framework [39, 40, 38] whereby treatment assignment (Z_i) is viewed as a randomized instrument that encourages individuals to update their beliefs about diversity in the local police department (D_i).

This approach rests on two core assumptions: 1) the instrument, Z, has a causal effect on beliefs, D_i; and 2) Z_i only affects outcomes, Y_i, through the path $Z_i \rightarrow D_i \rightarrow Y_i$ (exclusion restriction assumption). Under these assumptions (the first is empirically testable), we can estimate the (local) Average Treatment Effect (LATE) of D_i on Y_i. Specifically, we leverage the random assignment of information about police diversity (Z_i) to quantify the causal effect of changing beliefs about police diversity (D_i) on outcomes (Y_i) with instrumental variables regression using Two-Stage Least Squares (2SLS). The 2SLS estimator,

$$\widehat{\beta}_{IV} = \frac{\widehat{\mathrm{Cov}}(Y_i, Z_i)}{\widehat{\mathrm{Cov}}(D_i, Z_i)} = \frac{\widehat{\mathrm{Cov}}(Y_i, Z_i)/\widehat{\mathrm{Var}}(Z_i)}{\widehat{\mathrm{Cov}}(D_i, Z_i)//\widehat{\mathrm{Var}}(Z_i)}$$

is the ratio of the "reduced-form" effect of treatment assignment (the "instrument" Z_i) on a given outcome Y_i, and the "first-stage" effect on their beliefs about police diversity, D_i. Given random assignment of the instrument, $\widehat{\beta}_{IV}$ is consistent for the causal effect of belief updating on outcomes, provided the exclusion restriction assumption holds and the first-stage effect is non-zero.

We measured belief updating with the following post-treatment question: "To what extent do you agree or disagree with the following statement: "The Yonkers Police Department (YPD) adequately reflects the diversity of the community it serves." Responses were recorded on a 7-point scale (reverse coded) from "Strongly agree" (1) to "Strongly disagree" (7) with a neutral midpoint (4). Among those assigned to the information condition (treatment) 71% provided a response above the neutral midpoint, compared with 52% in the no information condition (control). The average was 4.54 scale points (SE = 0.09) in the control group and 5.27 (SE = 0.09) in the treatment group. Therefore, the estimated "first-stage effect" from OLS regression of belief updating on treatment assignment is 0.73 scale points ($\mathrm{SE} = 0.13, P < 0.01$). The estimated F-statistic from this regression is 32.44 ($P < 0.01$), well above the recommended threshold of 10 used to distinguish "weak" from "acceptable" instruments in applied work [41, 42]. This provides clear evidence that treatment assignment is a strong instrument for belief updating in this setting.

Table S10 compares estimates of the reduced-form effects of treatment assignment on outcomes with the 2SLS estimates for the (local) causal effects of belief updating on outcomes. As these results demonstrate, the causal effects of belief updating on support for tie-breaking, voting for police diversification, and trust and confidence in the police were all statistically distinguishable from zero and in the expected direction. Moreover, the 2SLS estimates are stronger in magnitude for these outcomes. For example, the reduced form estimate of the effect of information on the voting outcome is approximately 7 percentage points, whereas the 2SLS estimate of the effect of belief updating is approximately 10 percentage points.

	Reduced-form	2SLS
Support for affirmative action in recruitment and hiring	0.03 (0.07)	0.04 (0.10)
Support for tie-breaking in favor of minority applicants	0.22 (0.07)*	0.30 (0.09)*
Voted for police department diversification	0.07 (0.03)*	0.10 (0.04)*
Donation to Black police officer association	-0.62 (1.58)	-0.84 (2.17)
Trust and confidence in the police department	-0.13 (0.04)*	-0.17 (0.06)*
Willingness to cooperate with police officers	0.02 (0.04)	0.03 (0.06)

Table S10: Estimates from reduced-form regressions and instrumental variables regressions in municipal sample. The first column of results shows point estimates for the ATEs from OLS regressions of the outcome on treatment assignment, with robust standard errors in parentheses. The second column of results shows point estimates (standard errors) for the Local Average Treatment Effects (LATEs) from instrumental variables regressions using two-stage least squares (2SLS). $^*P < 0.05$

We also explored whether information salience might instead provide a better explanation for the effects in the information experiments. That is – does receiving novel information about police diversity simply increase the perceived importance of minority representation in the minds of respondents? To do so, we leverage two additional post-treatment questions (presented in randomized order): 1) "In your view, how important is it that police officers closely resemble the communities they serve in terms of gender?"; 2) "In your view, how important is it that police officers closely resemble the communities they serve in terms of race/ethnicity?" Responses were captured using a 5-point scale: "Not at all important" (1), "Slightly important" (2), "Moderately important" (3), "Very important" (4), or "Extremely important" (5). We combine these items into a single index ($\alpha = 0.73$, range 1-5) to obtain a measure of perceived importance of minority representation in the police.

If providing novel information about police diversity changed attitudes largely because it caused respondents to attach more importance to the issue of minority representation then we would except to see positive effects on these measures. That is, we should expect to observe a strong positive "first stage" effect from a regression of this measure on treatment assignment. Unlike our measure of belief updating, however, we do not find strong evidence in support of this. The average was 3.30 (SE = 0.07) in the information condition (treatment) and 3.21 (SE = 0.06) in the no information condition (control). The estimated first stage effect from the regression on treatment assignment was 0.09 scale points on a 5-point scale (SE = 0.09, P = 0.30), with an F-statistic of 1.08 (P = 0.30). The next section provides additional evidence (from the national sample) that treatments which emphasize the potential benefits of police diversification for under-represented minority groups only cause attitude change when paired with information.

Overall, these results are consistent with the idea that information provision increased support for diversification (and reduced trust) via belief updating. By comparison, we find weak evidence for the salience mechanism, i.e. that information about the lack of minority representation in the police department caused respondents to attach more importance to the issue of minority representation. We interpret these analyses as demonstrating that, on average, respondents attach a moderate amount of importance to police diversity regardless of the information they have available. Exposure to factual information about the lack of diversity in their local police force did not meaningfully increase issue salience, but did cause respondents to revise their beliefs about police diversity. This decreased their trust and confidence in the police, and increased their support for diversification.

S2.1.5 Average treatment effects of additional treatment arms in national sample

As described in Section S1.3, the information provision experiment fielded on the national sample consisted of four treatment arms: 1) no information (control); 2) information about police diversity only ("Info treatment"); 3) information about a recent *Science* publication [21] describing the potential benefits of police diversification for minority residents ("*Science* treatment"); 4) both information about police diversity and the *Science* article ("Info + *Science*"). Thus far, we have focused attention on the estimated ATEs of the Info condition (relative to control) for comparison with the municipal sample (which only assigned these two conditions).

Table S11 provides estimates for all three ATEs in the national sample (each relative to control). To facilitate comparisons, all estimates are standardized using using Glass's Δ, which scales outcomes by the standard deviation in the control group [28, 29]. These results show that exposure to high-quality research demonstrating the potential benefits of police diversification for minority groups did not, on its own, cause attitude change. Instead, we find that the estimated ATE of information provision is statistically indistinguishable from the effect of information provision *and* relevant research, again demonstrating the powerful effects of exposure to information about police diversity on its own.

	Information	**Science**	**Information + Science**
Support for affirmative action in recruitment and hiring	0.00 (0.06)	-0.06 (0.06)	-0.07 (0.06)
Support for tie-breaking in favor of minority applicants	0.17 (0.06)*	0.01 (0.06)	0.17 (0.07)*

Table S11: Estimated treatment effects of additional treatment arms in national sample. Point estimates for the ATEs estimated using OLS regression of the outcome on treatment assignment, with robust standard errors in parentheses. All estimates are standardized using Glass's Δ, which scales outcomes by the standard deviation in the control group [28, 29].
$^{*}P < 0.05$

S2.1.6 Average treatment effects on additional outcomes in municipal sample

Outcome	**Estimate**
Rank ordering of diversification policy	0.16 (0.08)
Stated support for diversification policy	0.13 (0.14)
Willigness to consider policing career	0.02 (0.07)
Importance of police officer diversity:	
Race/ethnicity	0.02 (0.10)
Gender	0.17 (0.10)
Beliefs about diversity among US police in general:	
White officer share	4.38 (1.20)*
Black officer share	-2.07 (0.56)*
Hispanic/Latino officer share	-1.09 (0.62)
Asian officer share	-1.21 (0.39)*
Male officer share	2.40 (1.10)*
Female officer share	-2.40 (1.10)*

Table S12: Estimated treatment effects on additional outcomes in municipal sample. Point estimates for the ATEs estimated using OLS regression of the outcome on treatment assignment, with robust standard errors in parentheses. $^*P < 0.05$

Description of additional outcome measures:

- **Stated support for diversification policy**. Support for diversification policy described in Section S2.1.2. Measured in both the baseline and followup survey.
- **Rank ordering of diversification policy**. Relative importance of diversification policy described in Section S2.1.2. Measured in both the baseline and followup survey.
- **Willingness to consider policing career (2-item index)**. Respondents were provided with updated data on the salary and benefits for civil service occupations as a police officer, firefighter, and school teacher. For each occupation, they were asked how likely they would be to consider a career in this occupation, and how likely they would be to encourage a close friend or family member to consider a career in each (see Fig. S33). Each question was presented in random order and recorded using a 7 point scale from "Extremely unlikely" to "Extremely likely" with a neutral midpoint. A 2-item index was created using responses to the police officer questions. Only measured in the followup survey.
- **Importance of racial diversity among police**. Responses to the question "In your view, how important is it that police officers closely resemble the communities they serve in terms of race/ethnicity?" recorded on a 5 point scale: 1 = "Not at all important", 2 = "Slightly important", 3 = "Moderately important", 4 = "Very important", 5 = "Extremely important".
- **Importance of gender diversity among police**. Responses to the question "In your view, how important is it that police officers closely resemble the communities they serve in terms of gender?" recorded on a 5 point scale: 1 = "Not at all important", 2 = "Slightly important", 3 = "Moderately important", 4 = "Very important", 5 = "Extremely important".

- **Beliefs about diversity among US police in general**. Responses to the same questions that were used to capture pre-treatment beliefs about police officer diversity in the national sample (see Fig. S2).

If **you were in a position to start a new career**, how likely would you be to consider each of the following civil service occupations listed below?

The entry-level salaries and median pay for Yonkers, NY are provided below, based on the most recent government data from 2020. Benefits for each occupation are comparable, and include health insurance during your entire career and eligibility for a New York State Pension after 20 years of service.

Teacher at Yonkers Public Schools. Entry-level salary between $68,662 and $86,148 depending on education level (Bachelors, Master, or PhD). The median total pay among all Yonkers public school teachers was $133,425 in 2020, and 75% earned $137,000 or higher.

Firefighter at Yonkers Fire Department. Entry-level salary of $71,605 (high school diploma or GED required). The median total pay among all Yonkers firefighters was $142,542 in 2020, and 75% earned $129,000 or higher.

Police officer at Yonkers Police Department. Entry-level salary of $72,233 (high school diploma or GED required). The median total pay among all Yonkers police officers was $132,383 in 2020, and 75% earned $111,000 or higher.

Figure S33: Willingness to consider a career as a police officer question

S2.1.7 Causal attributions for lack of police officer diversity

In this section we report descriptive evidence regarding Yonkers residents' belief that specific factors explain disparities in minority representation in U.S. police forces. At the end of our second municipal survey, we explained to all respondents that there are many police stations across the U.S. that underrepresent the minority communities they serve. After doing so, we asked respondents to express the extent to which they believe that a list of four factors explains disparities in representation.

Factors included: lack of demand on behalf of police forces to recruit minorities, lack of interested amongst minorities in joining police forces, lack of qualifications to serve as police officers amongst minority residents, and a lack of supporting environment in police departments. Generally, our descriptive results reported in Figure S34, suggest that there is quite a bit of variation across respondents with regards to causal attributions for lack of police officer diversity. However, it appears that attributions relating to minority “lack of qualifications" are widely dismissed, and perceptions that police departments may not provide minority officers with a supportive working environment are broadly endorsed.

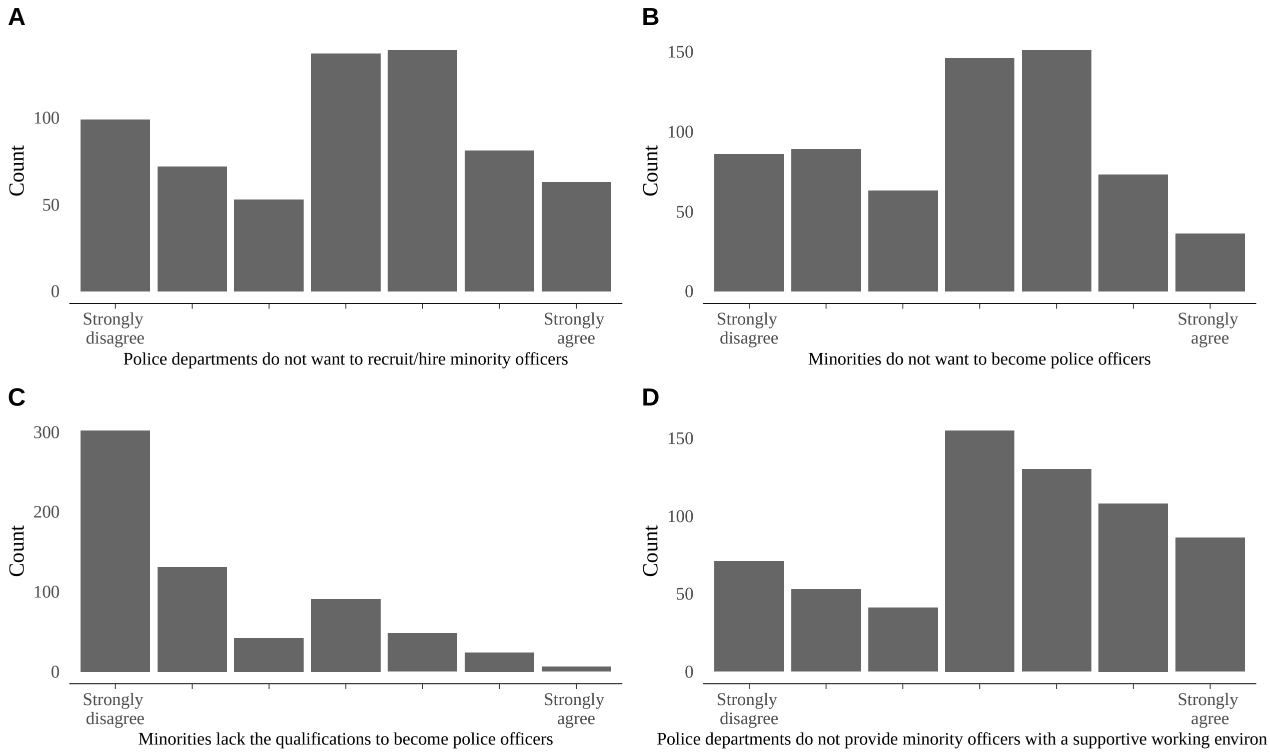

Figure S34: Causal Attribution of disparities in minority representation in U.S. police forces. Panels A-D represent the distribution of Yonkers residents' beliefs regarding possible factors that explain lack of diversity and minority representation in American police forces. These factors relate to lack of demand on behalf of police forces to recruit minorities (A), lack of interested amongst minorities in joining police forces (B), lack of qualifications to serve as police officers amongst minority residents (C), and a lack of supporting environment in police departments.

S2.1.8 Correlates of misperceptions about police diversity

In this section we examine demographic correlates of misperceptions regarding police diversity. To do so, we regressed variables capturing respondents misperceptions regarding the share of White, Male, Black, Latino, and Asian officers, over five different individual level binary indicators taking a value of one if a given respondent is: White, Male, Republican, Democrat, and college educated. In Figure S35, we report conditional correlations, for both our municipal and national samples.

Our exploration of conditional correlations in Figure S35 yields limited consistent evidence regarding systematic subgroup variation in misperception with and across samples. In our National sample, it appears that White survey respondents over-estimate the share of White and Male officers, while underestimating the share of minorities in police forces nationwide. This

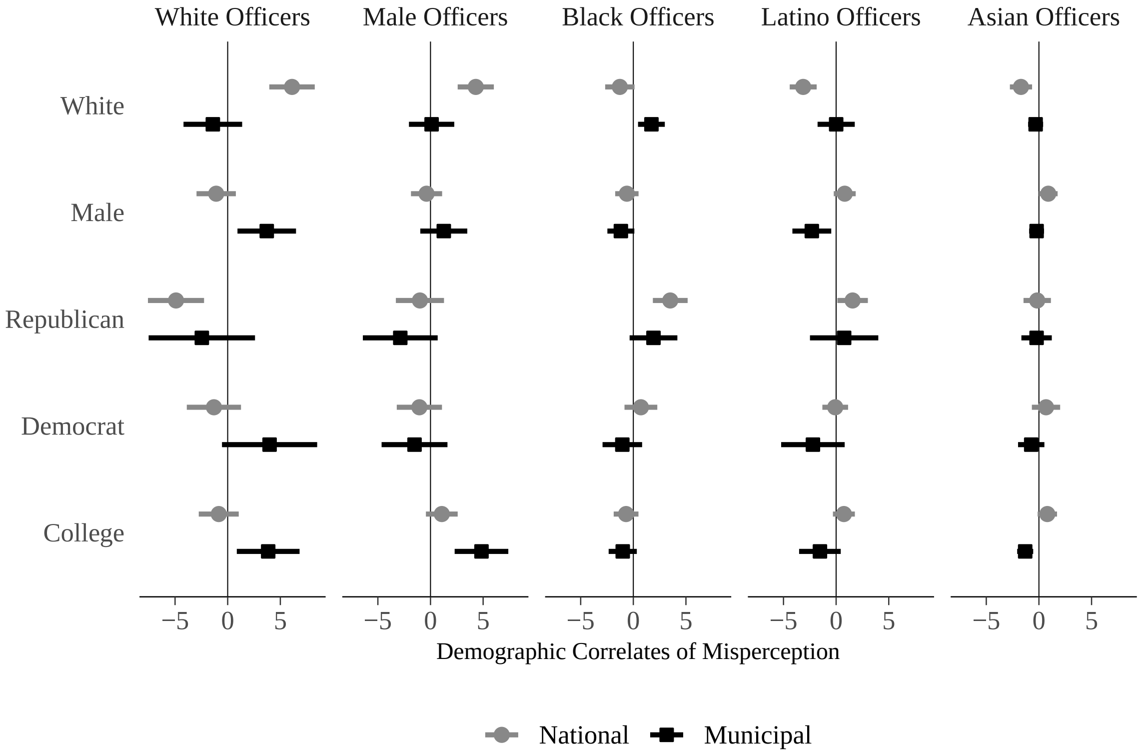

Figure S35: Demographic correlates of misperceptions about police officer diversity. Point estimates and 95% confidence intervals from OLS regressions with robust standard errors, estimating the conditional correlation of demographic variables, with respondents misperception of the share of White, Male, Black, Latino, and Asian officers.

result stands in contrast to results from our municipal sample in which race does not correlate with misperceptions, with the exception of misperceptions regarding the share of Black officers, which is positive, suggesting that White respondents over-estimate the share of Black officers in Yonkers.

Another variable that consistently correlates with misperceptions in the national sample is partisanship. Specifically, it appears that Republicans underestimate the share of white officers, and overestimate the share of Black and Latino officers in the police nationwide. In our municipal sample, a somewhat similar pattern emerges, as Republican Yonkers residents overestimate the share of Black officers. We do not find evidence for a precisely estimated correlation between respondents Democratic identification and misperceptions in the national and municipal sample.

When considering the conditional correlation of college education with misperceptions, we find that Yonkers residents who obtained a college degree overestimate the share of White and male officers, while underestimating the share of minority officers in YPD. In contrast, we find limited evidence for a precisely estimated conditional correlation in the national sample.

Finally, in the municipal sample, male respondents overestimate the share of White officers, and underestimate the share of Black and Latino officers. In the national sample, we find limited evidence for a gender conditional correlation, though male respondents slightly overestimate the share of Asian and Latino officers.

S2.2 Police recruitment conjoint experiments

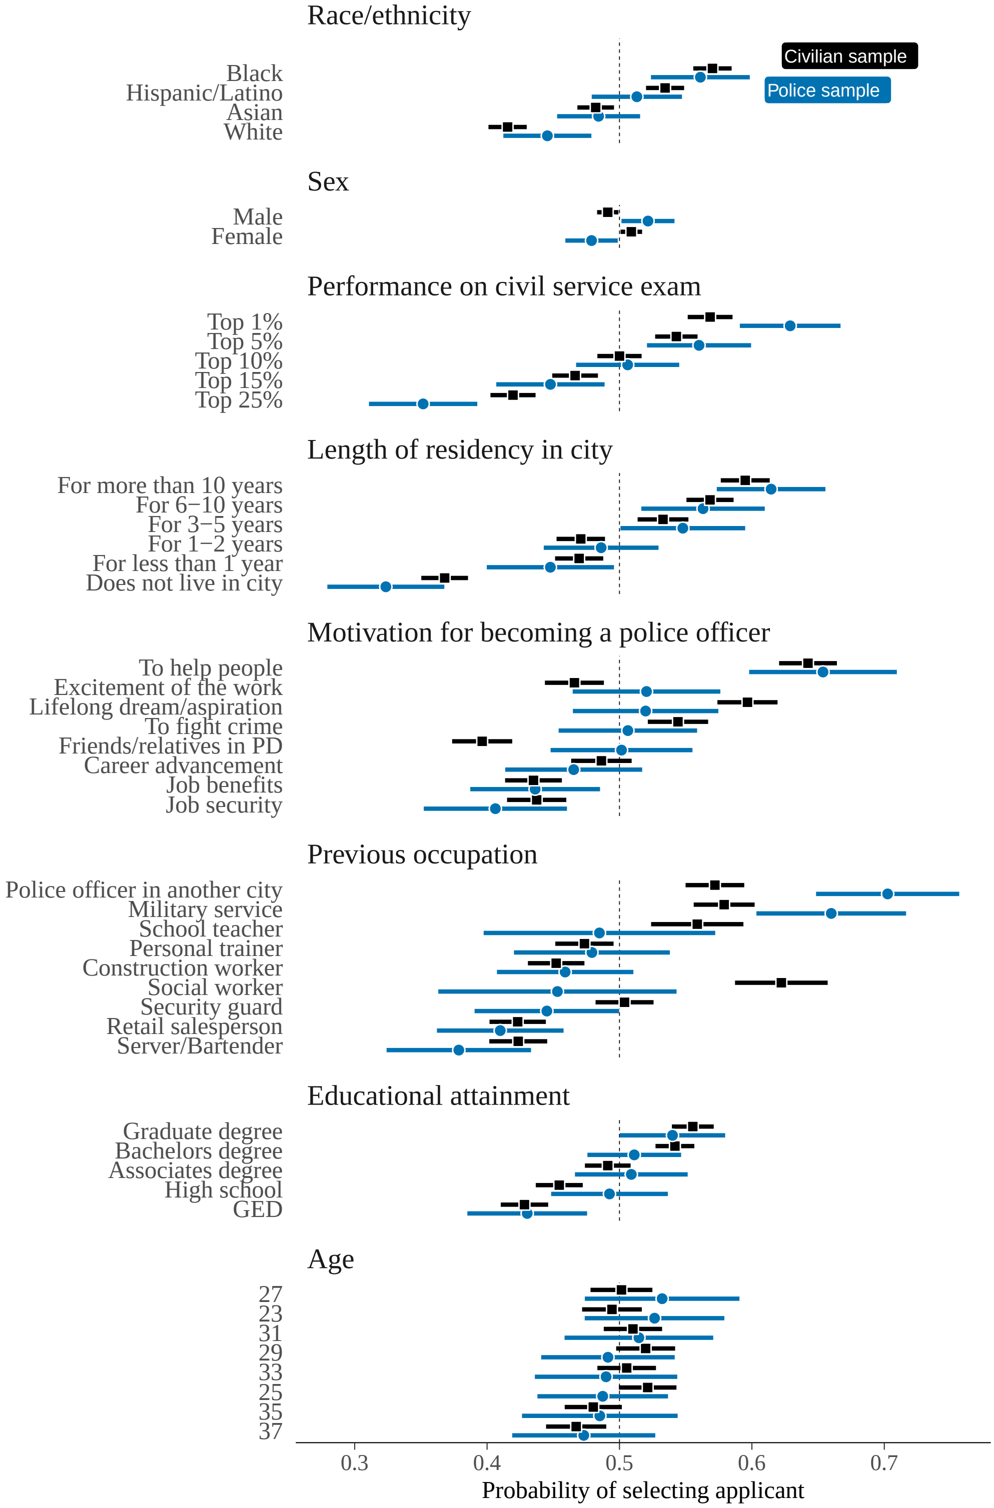

FIGURE S36: Estimated marginal means for all attributes in police recruitment conjoint by sample. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Estimates are adjusted to account for randomization constraints on the education and occupation attributes. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

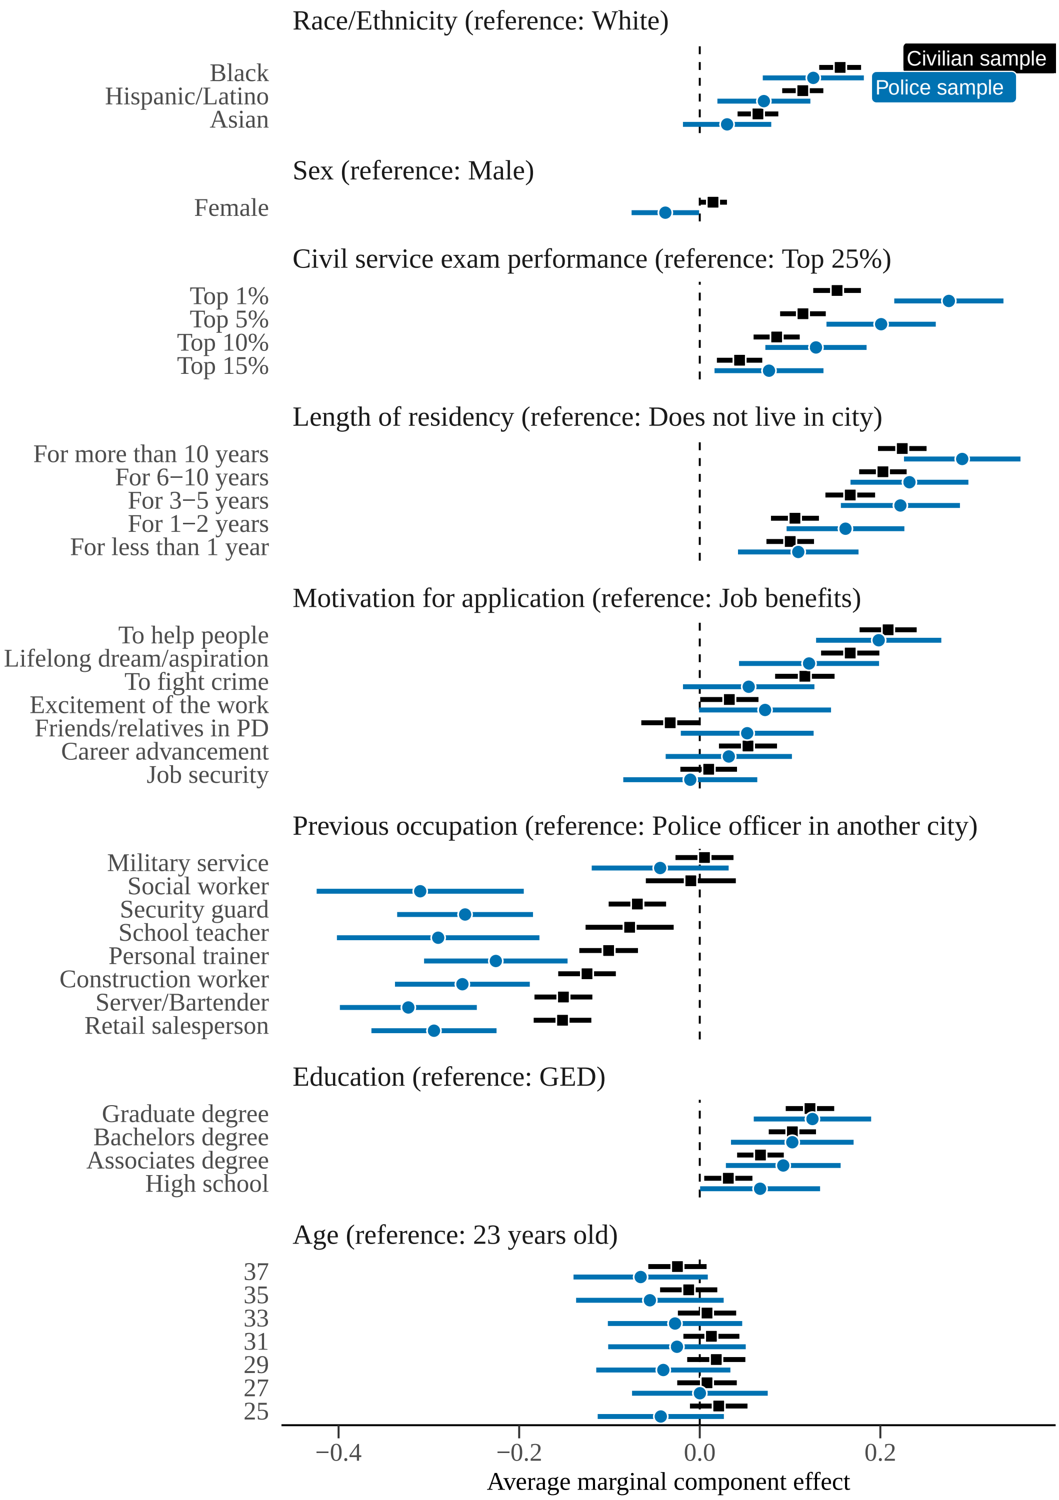

FIGURE S37: Estimated AMCEs for all attributes in police recruitment conjoint by sample. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Estimates are adjusted to account for randomization constraints on the education and occupation attributes. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

S2.2.1 Estimated marginal means and AMCEs on ordinal outcome measure

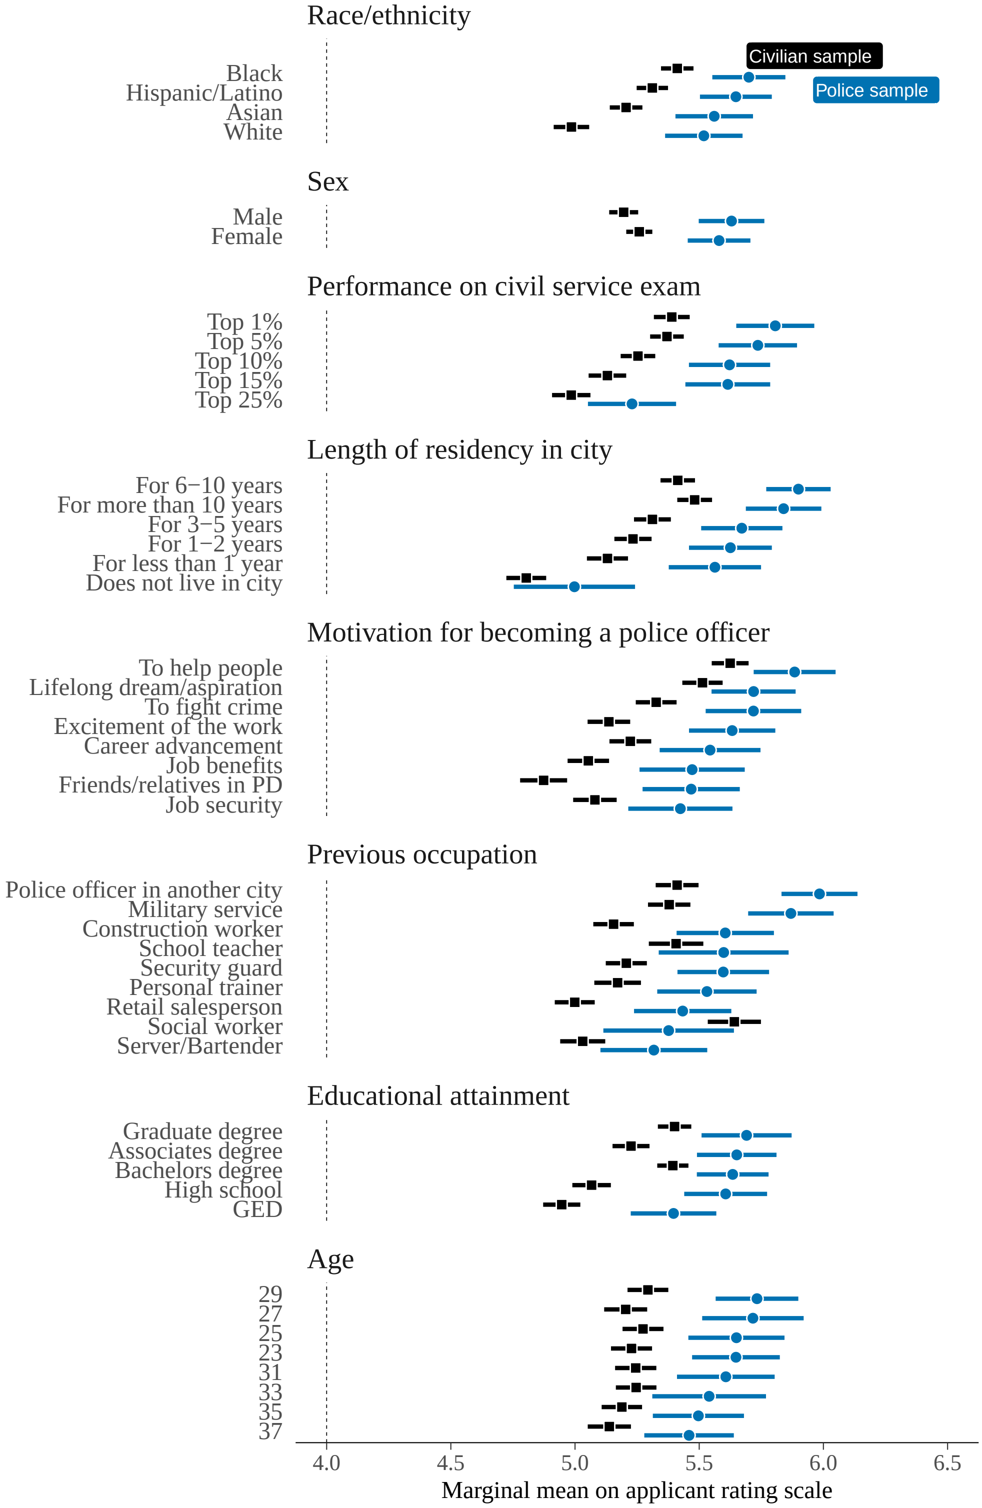

FIGURE S38: Estimated marginal means on ordinal outcome for all attributes in police recruitment conjoint by sample. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Estimates are adjusted to account for randomization constraints on the education and occupation attributes. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

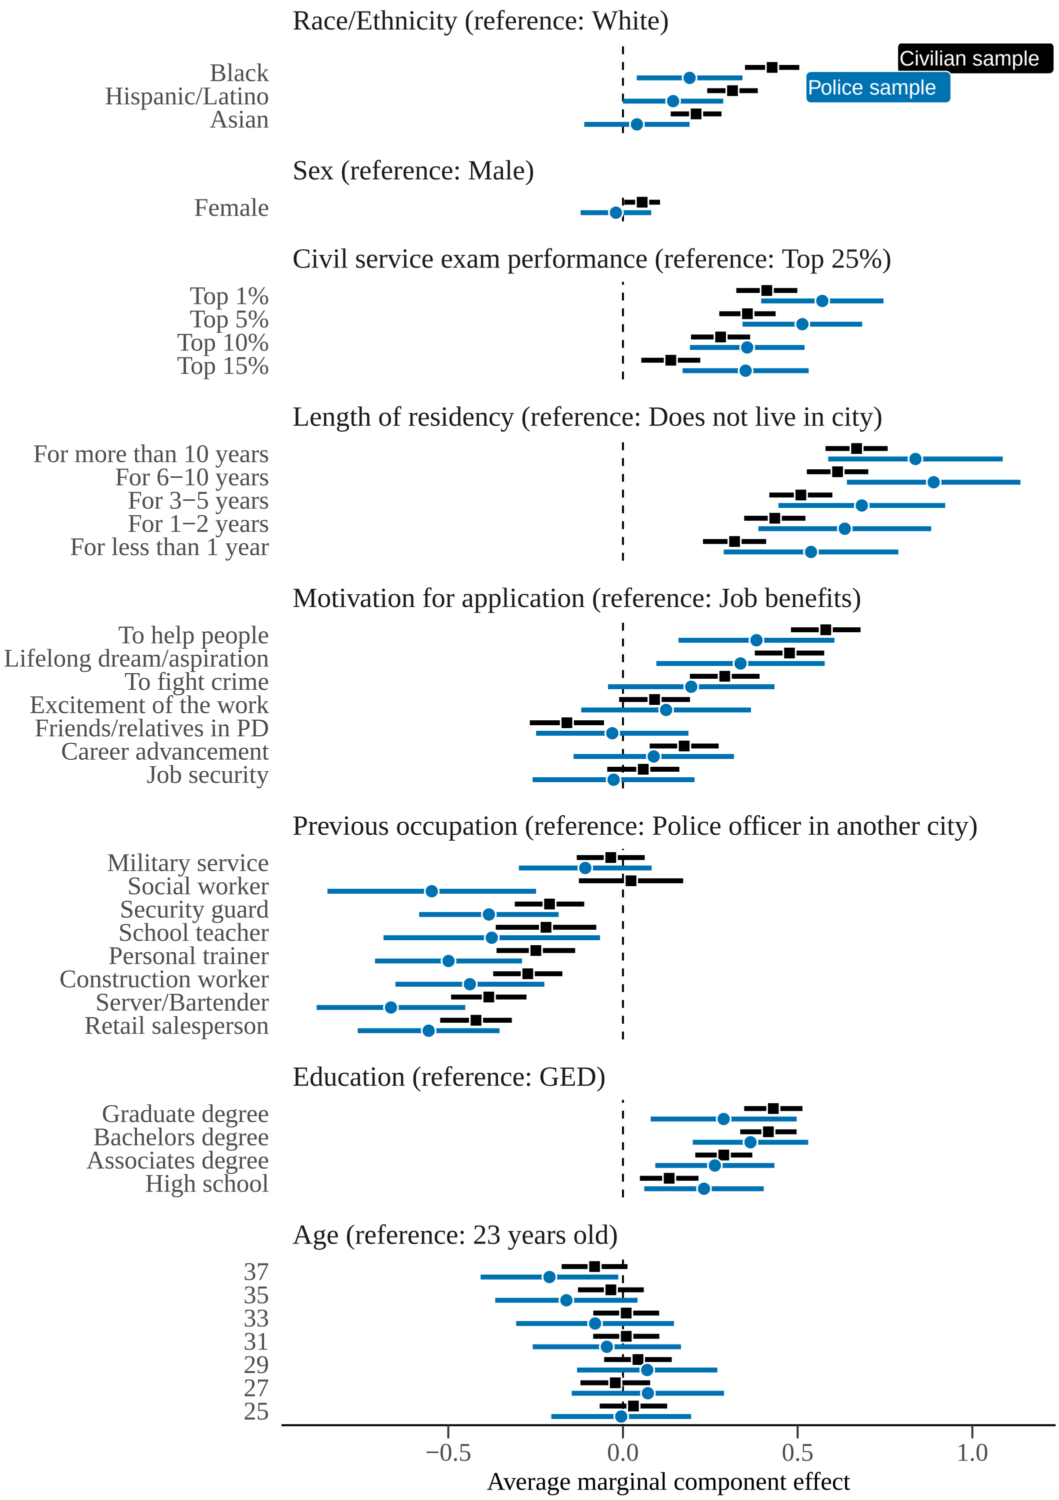

FIGURE S39: Estimated AMCEs on ordinal outcome for all attributes in police recruitment conjoint by sample. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Estimates are adjusted to account for randomization constraints on the education and occupation attributes. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

S2.2.2 Causal interactions for race/ethnicity, sex, and exam performance

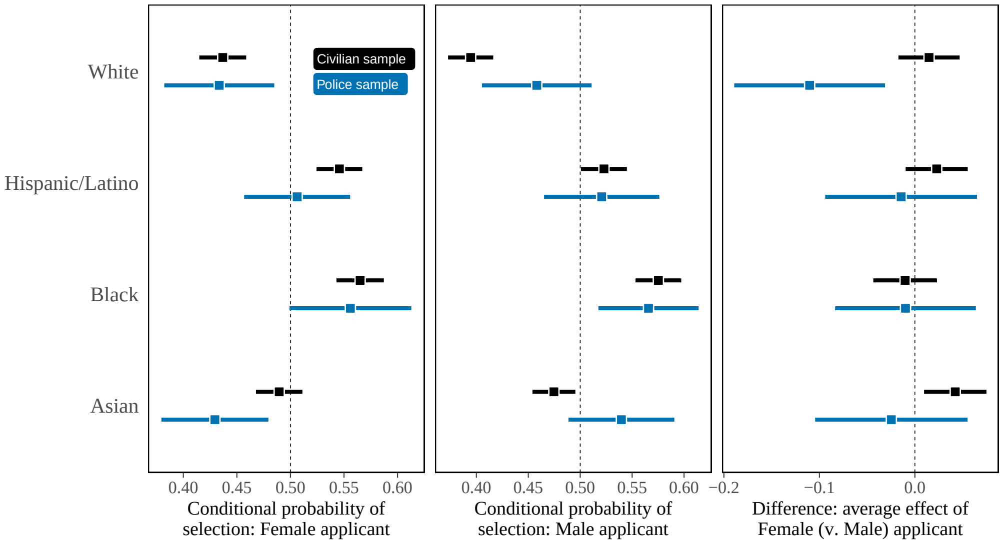

FIGURE S40: Estimated conditional marginal means for female applicants (left), male applicants (center), and the between sample differences (right). Differences capture the average causal effect of applicant sex (here: female v. male) on the probability of selection for each race/ethnicity category. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

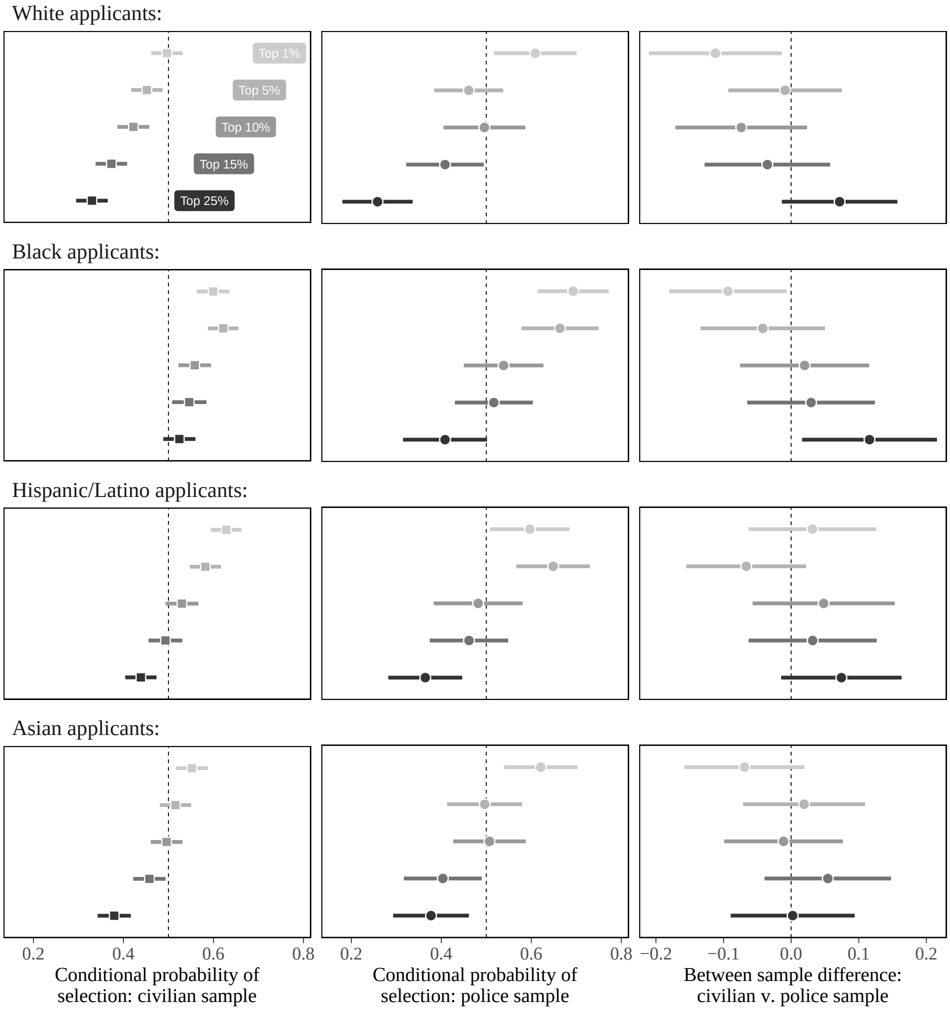

FIGURE S41: Estimated conditional marginal means by applicant race/ethnicity and civil service exam performance in civilian sample (left), police sample (center), and the differences (right) between samples. Positive (negative) differences indicate higher (lower) values in the civilian sample than the police sample. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

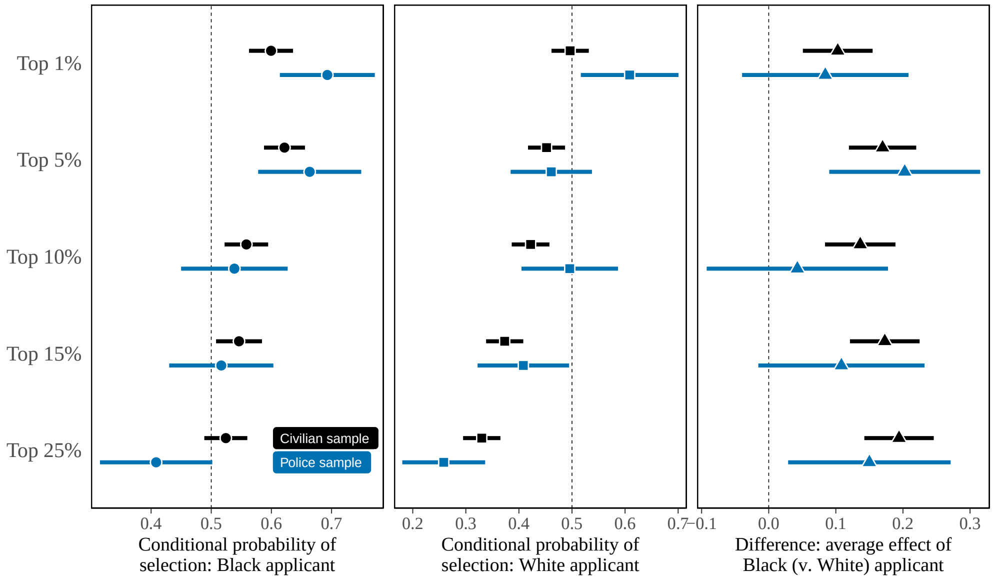

FIGURE S42: Estimated conditional marginal means for Black applicants (left), White applicants (center), and the differences (right) by civil service exam performance. Differences capture the average causal effect of applicant race/ethnicity (here: Black v. White) on the probability of selection at each level of the exam performance attribute. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

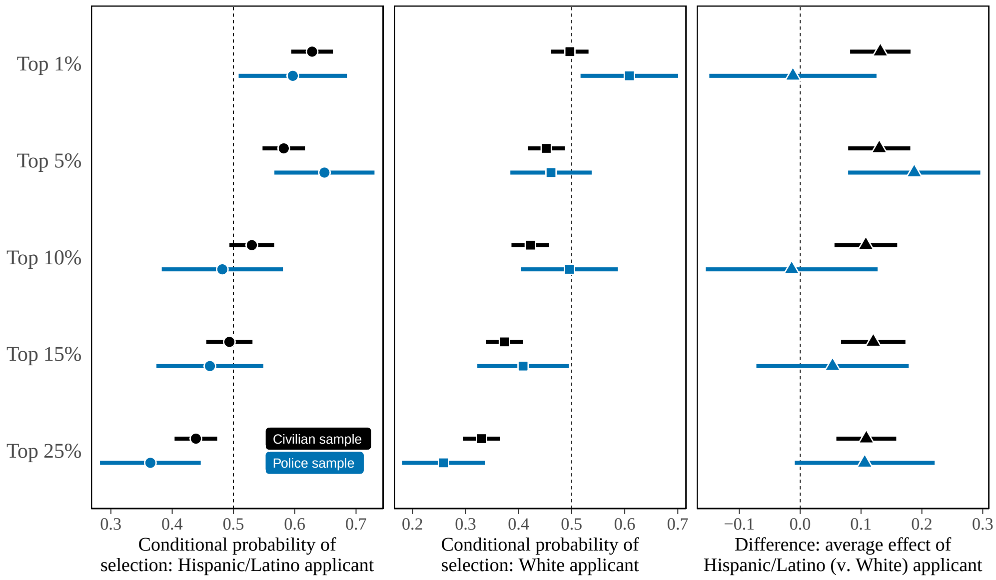

FIGURE S43: Estimated conditional marginal means for Hispanic/Latino applicants (left), White applicants (center), and the differences (right) by civil service exam performance. Differences capture the average causal effect of applicant race/ethnicity (here: Hispanic/Latino v. White) on the probability of selection at each level of the exam performance attribute. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

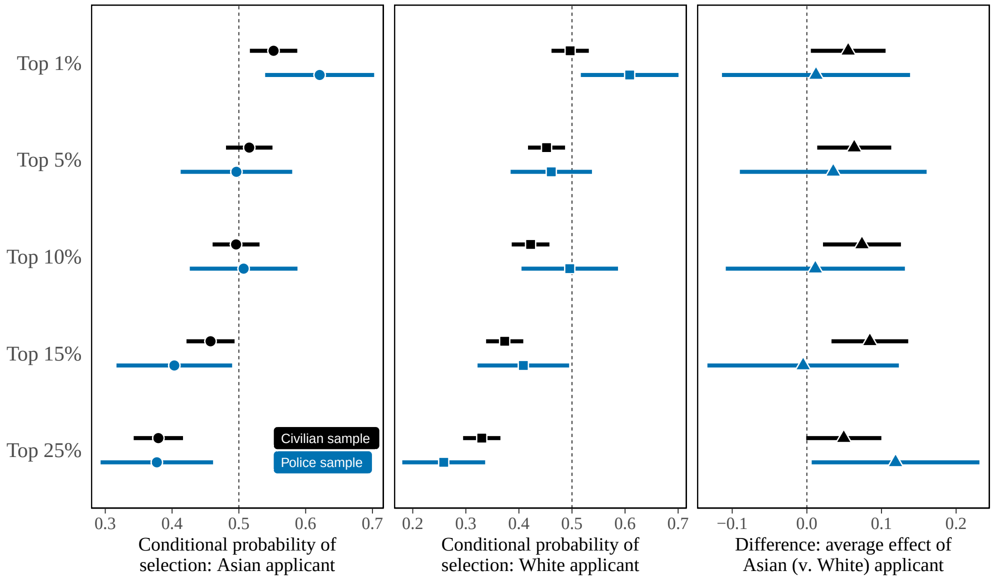

FIGURE S44: Estimated conditional marginal means for Asian applicants (left), White applicants (center), and the differences (right) by civil service exam performance. Differences capture the average causal effect of applicant race/ethnicity (here: Asian v. White) on the probability of selection at each level of the exam performance attribute. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

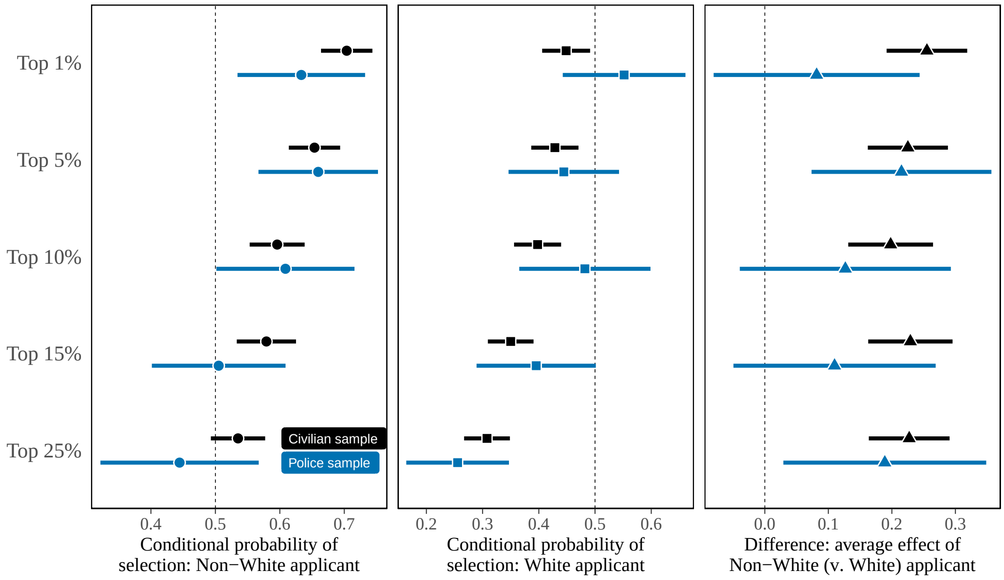

FIGURE S45: Estimated conditional marginal means for Non-White applicants (left), White applicants (center), and the differences (right) by civil service exam performance. The sample is restricted to the subset of randomized profiles that forced respondents to make pairwise comparisons between Non-White and White applicants (Civilian sample: 14,050 observations; Police sample: 2,440 observations). Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering.

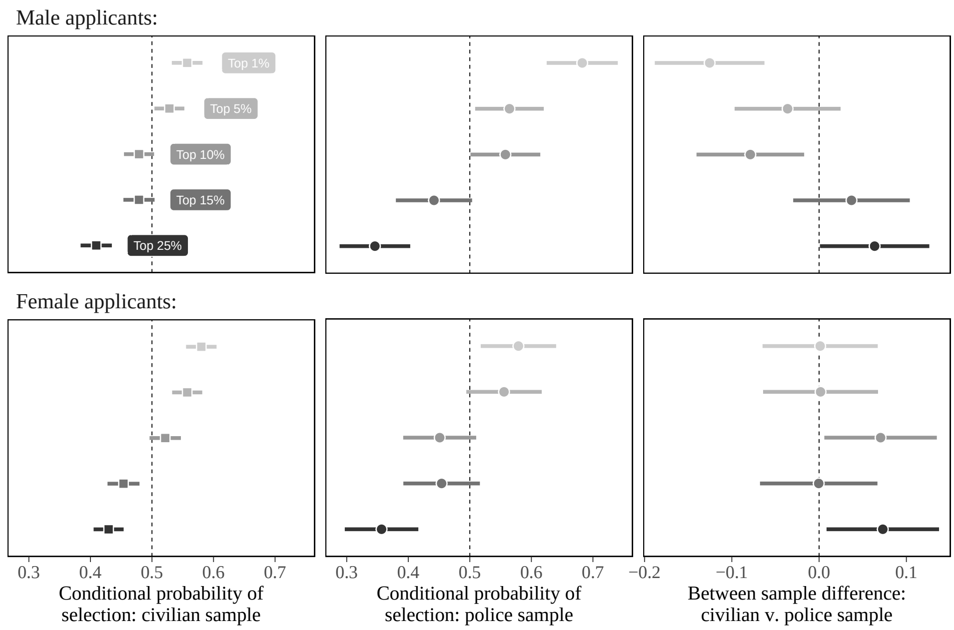

FIGURE S46: Estimated conditional marginal means by applicant sex and civil service exam performance in civilian sample (left), police sample (center), and the between sample differences (right). Positive (negative) differences indicate higher (lower) values in the civilian sample than the police sample. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

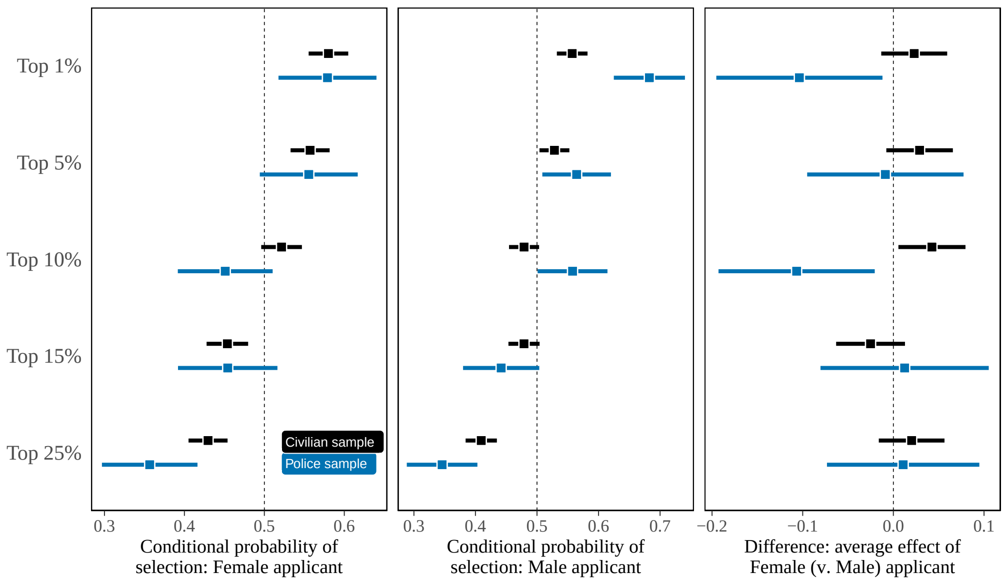

FIGURE S47: Estimated conditional marginal means for female applicants (left), male applicants (center), and the differences (right) by civil service exam performance. Differences capture the average causal effect of applicant sex (here: female v. male) on the probability of selection at each level of the exam performance attribute. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Civilian sample: municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations). Police sample: survey of Yonkers police officers fielded in June 2021 (N = 250 respondents x 5 pairings x 2 applicants per pair = 2,500 observations).

S2.2.3 Heterogeneity by respondent background characteristics

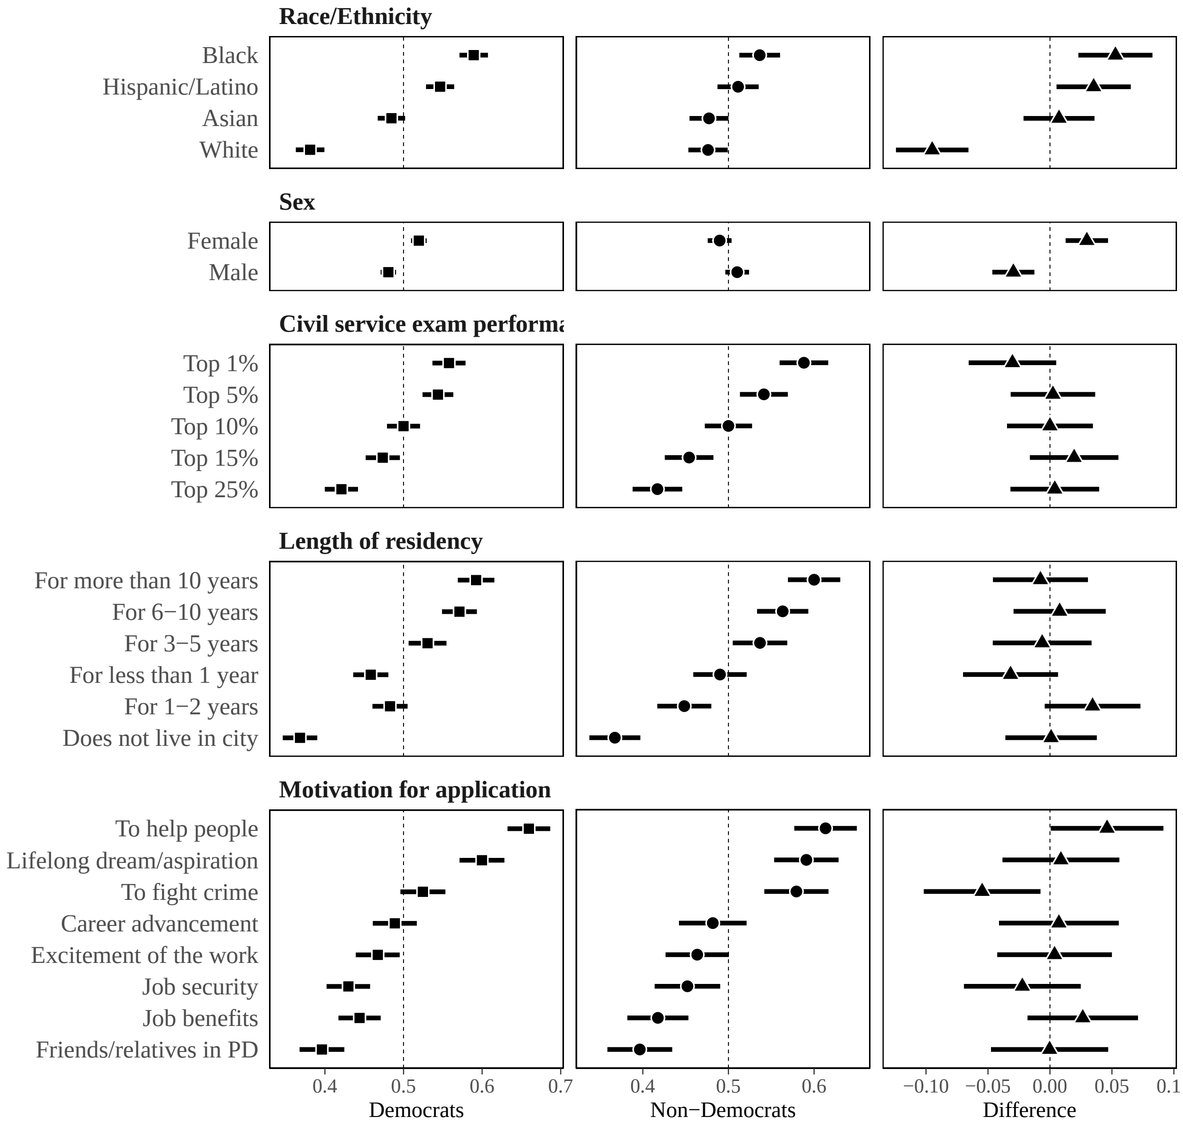

FIGURE S48: Estimated marginal means in police recruitment conjoint by partisanship. Sub-group estimates showing marginal means among Democrats (n = 913), non-Democrats (n = 500), and the differences. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

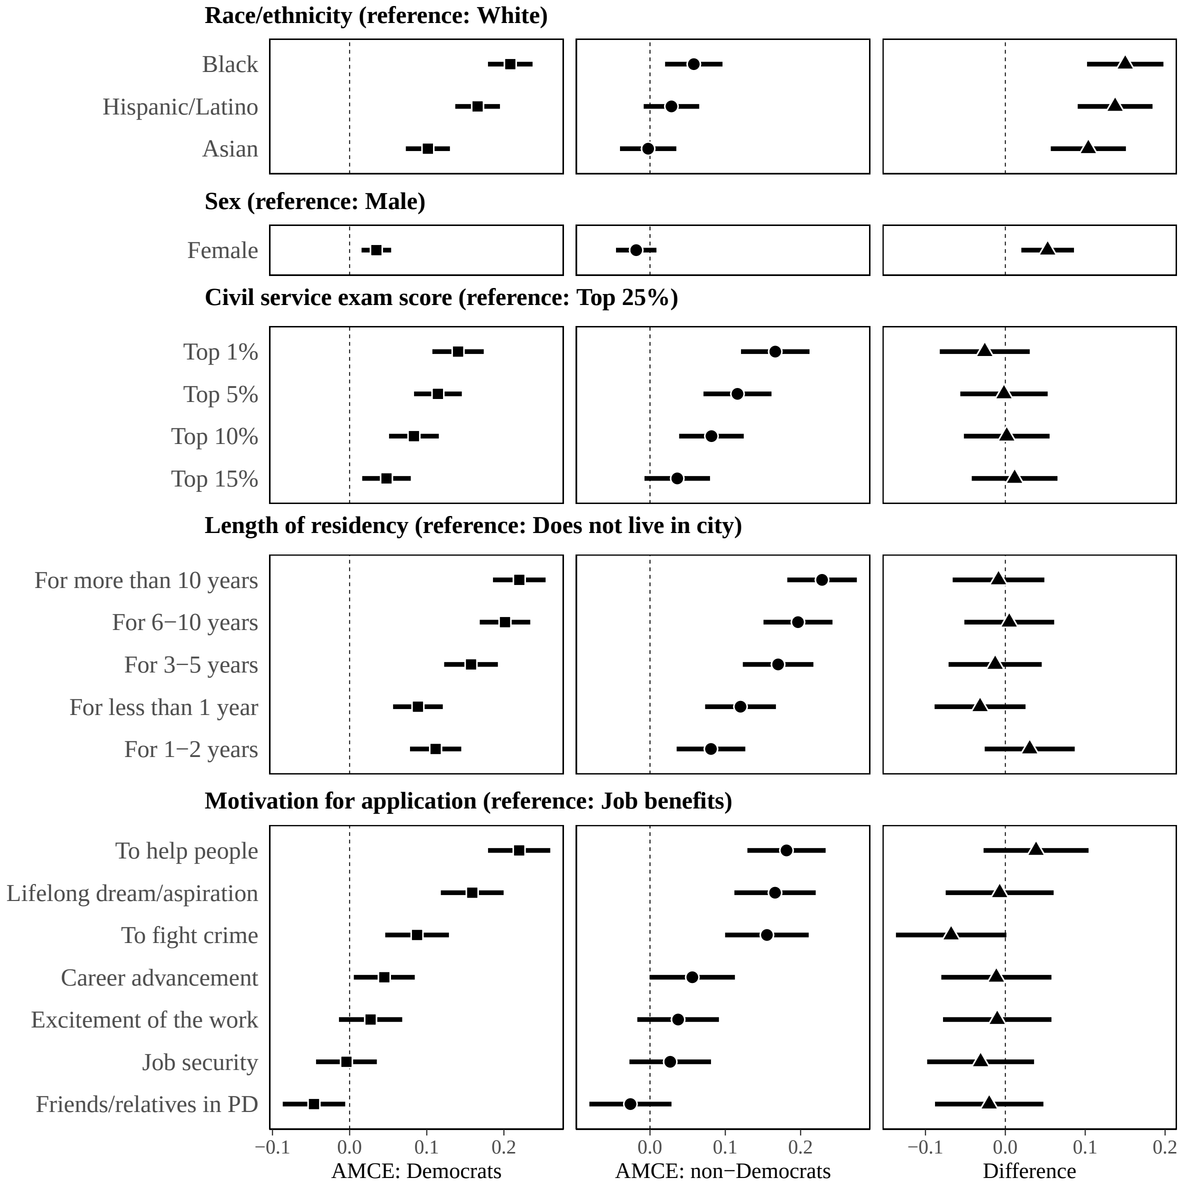

FIGURE S49: Estimated AMCEs in police recruitment conjoint by partisanship. Sub-group estimates showing AMCEs among Democrats (n = 913), non-Democrats (n = 500), and the differences. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

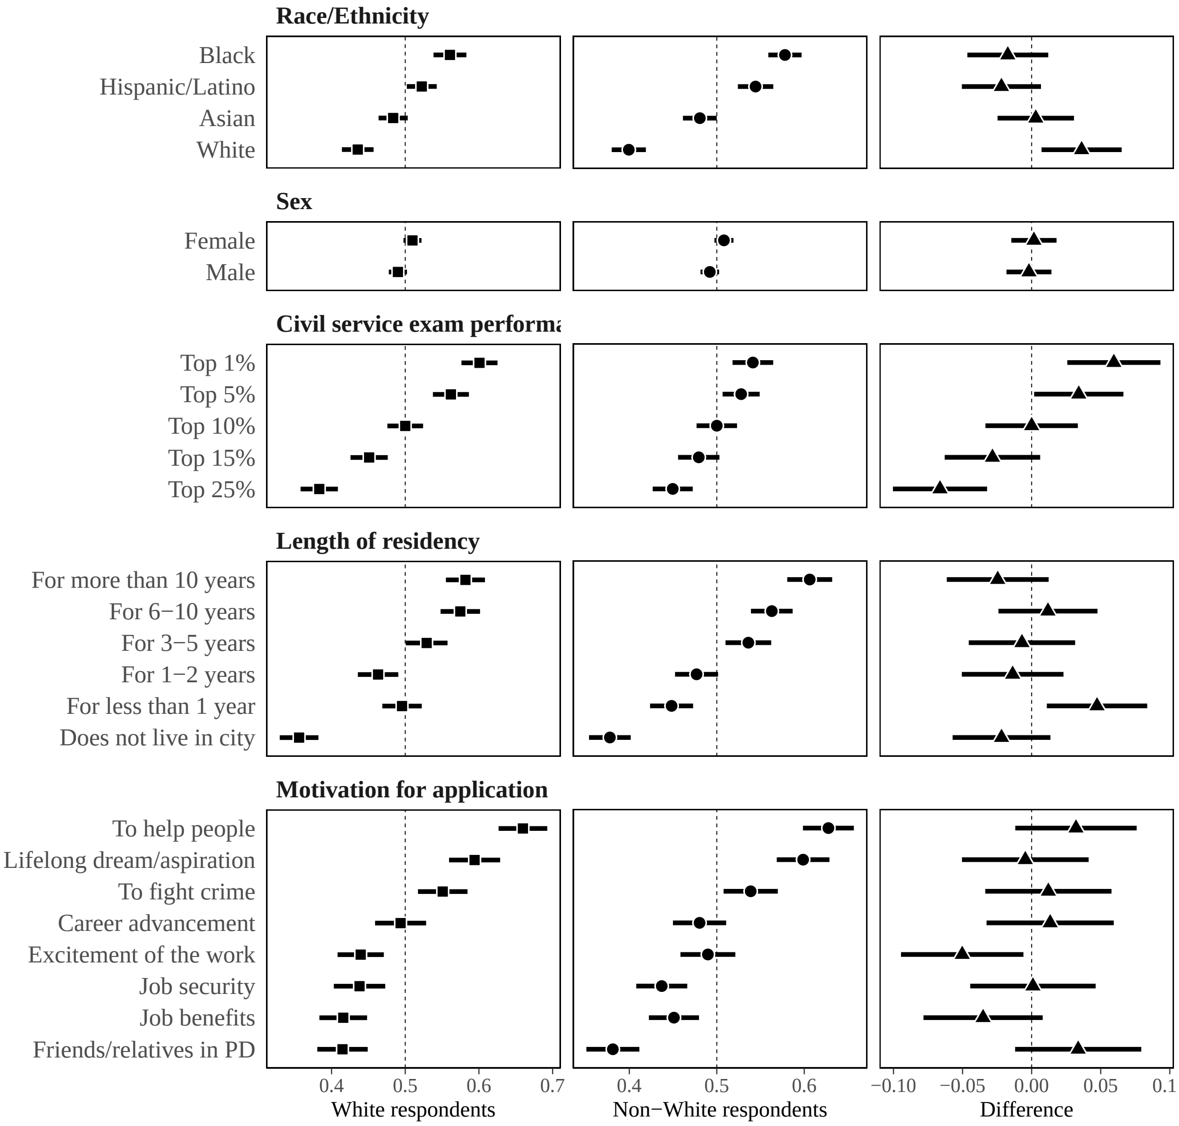

FIGURE S50: Estimated marginal means in police recruitment conjoint by race/ethnicity. Sub-group estimates showing marginal means among White respondents (n = 641), non-White respondents (n = 772), and the differences. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

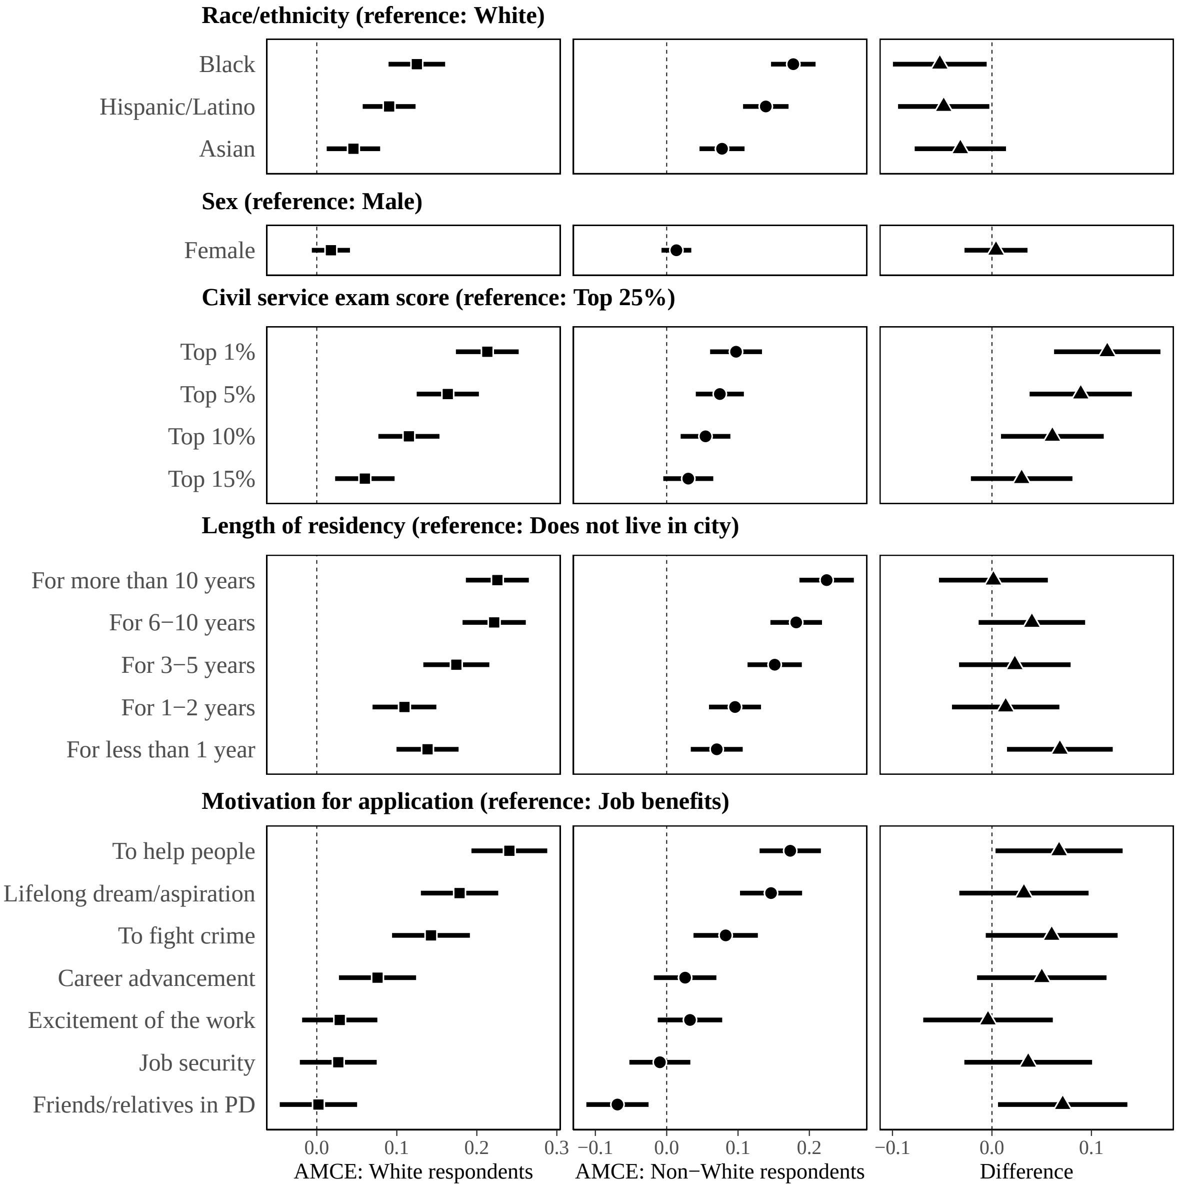

Figure S51: Estimated AMCEs in police recruitment conjoint by race/ethnicity. Sub-group estimates showing AMCEs among White respondents (n = 641), non-White respondents (n = 772), and the differences. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

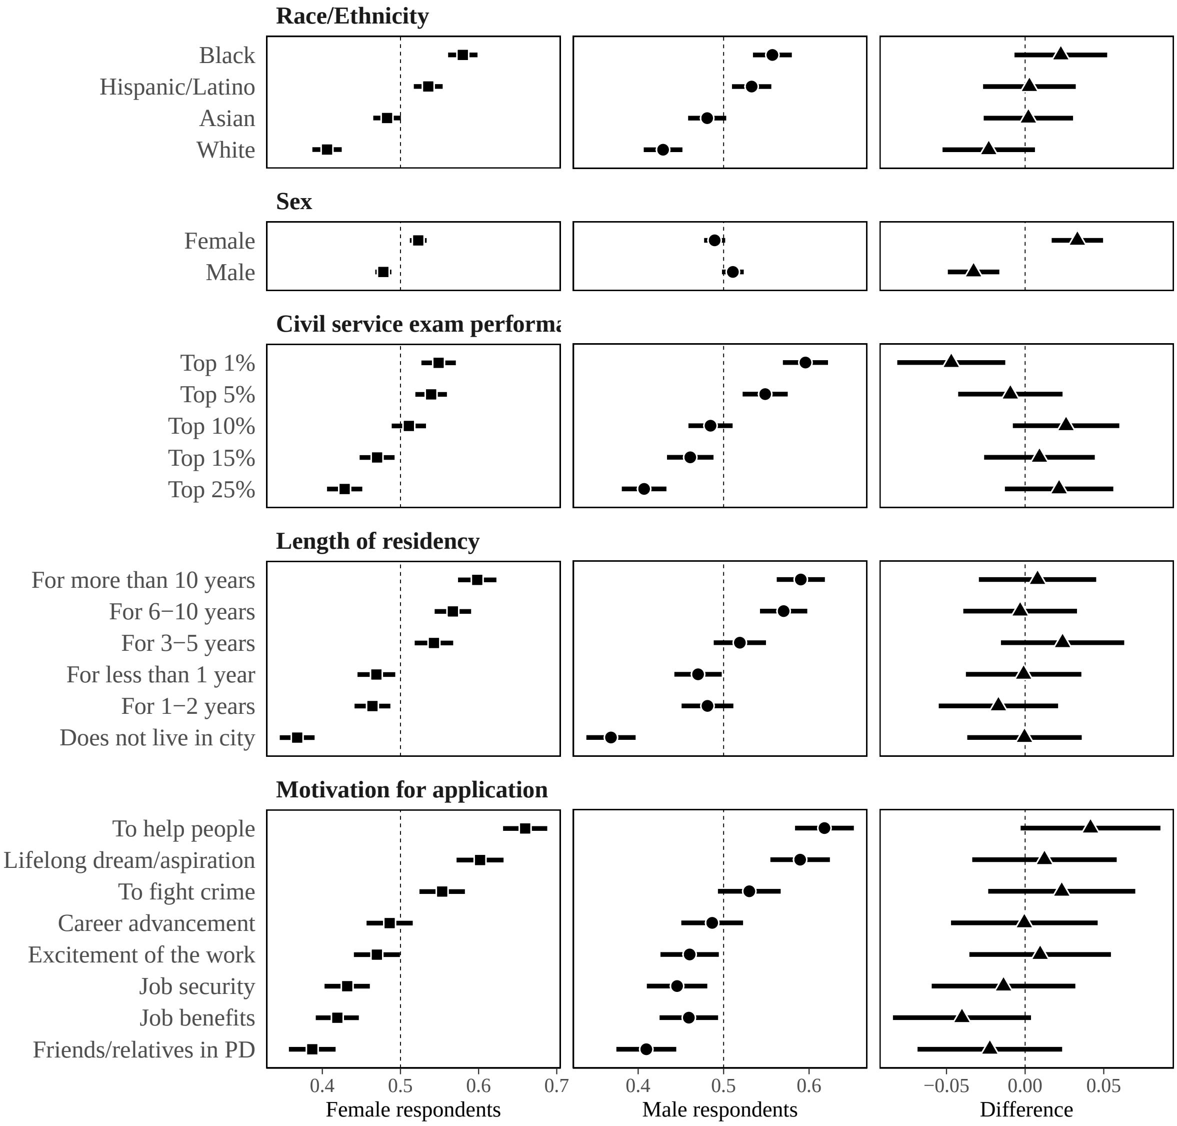

FIGURE S52: Estimated marginal means in police recruitment conjoint by sex. Sub-group estimates showing marginal means among male respondents (n = 576), female respondents (n = 837), and the differences. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

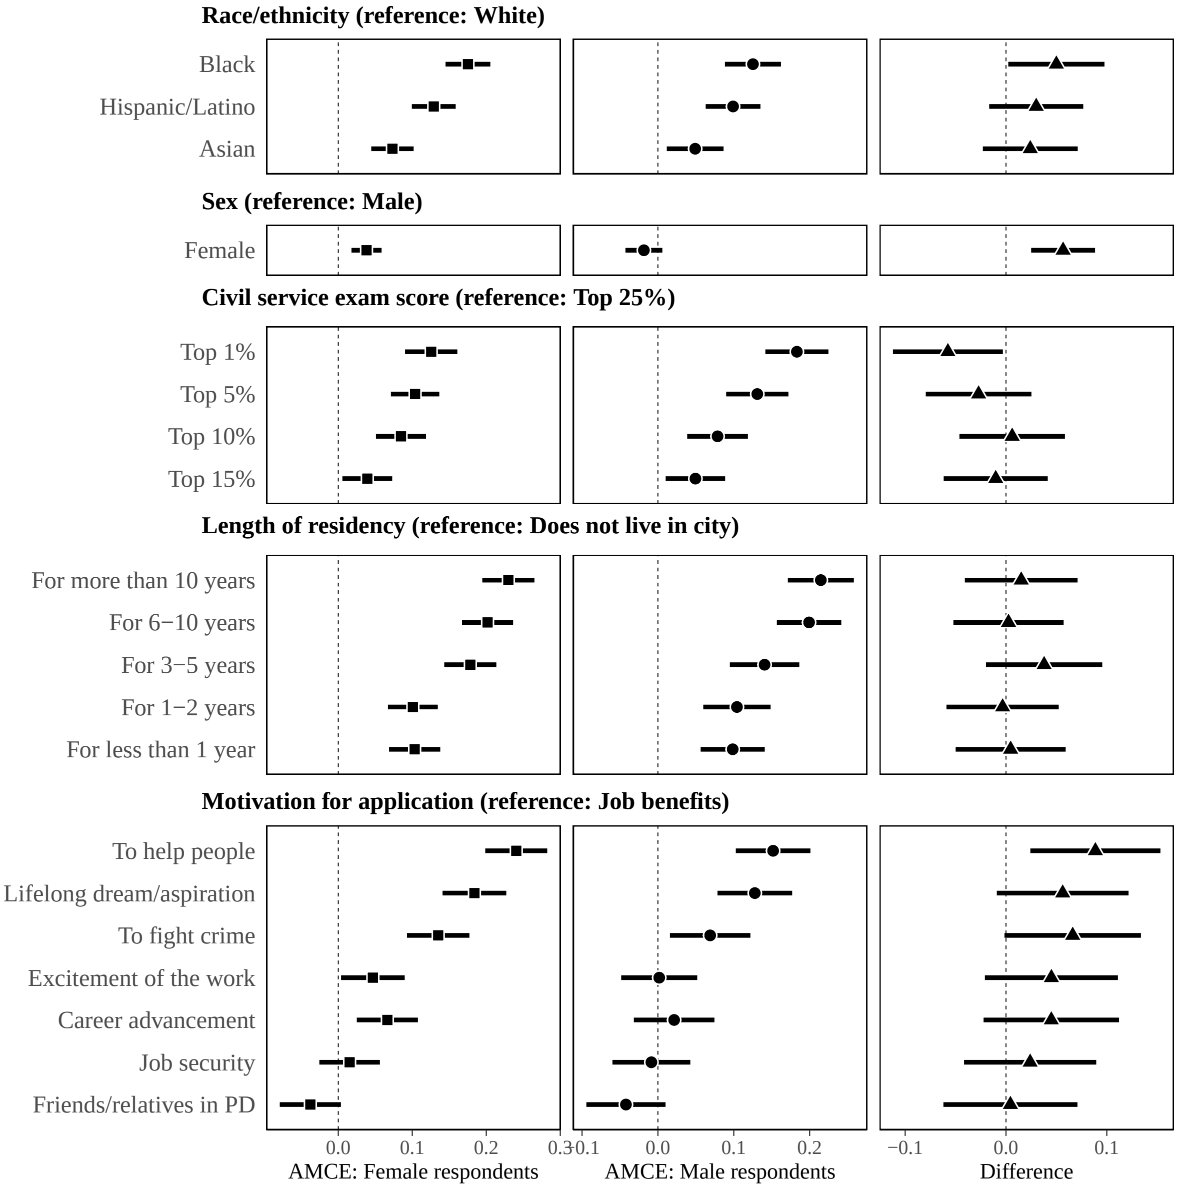

Figure S53: Estimated AMCEs in police recruitment conjoint by sex. Sub-group estimates showing AMCEs among male respondents (n = 576), female respondents (n = 837), and the differences. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

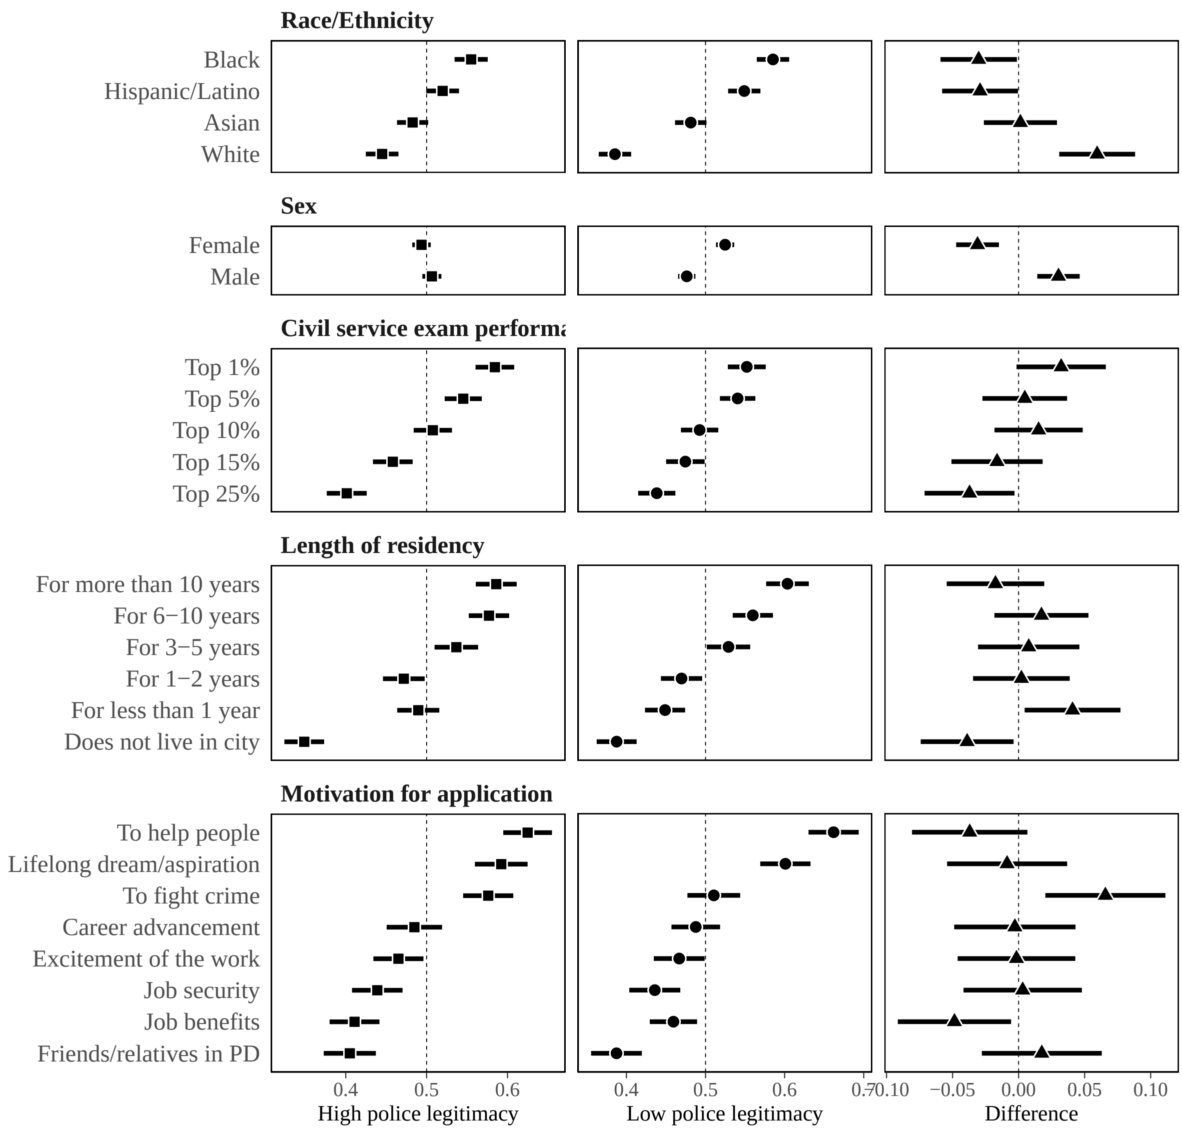

FIGURE S54: Estimated marginal means in police recruitment conjoint by police legitimacy. Sub-group classifications are based on scores derived from the 10-item index of police legitimacy, trust, and confidence. Police legitimacy is coded as high if a respondent scored higher than the median respondent on the index. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

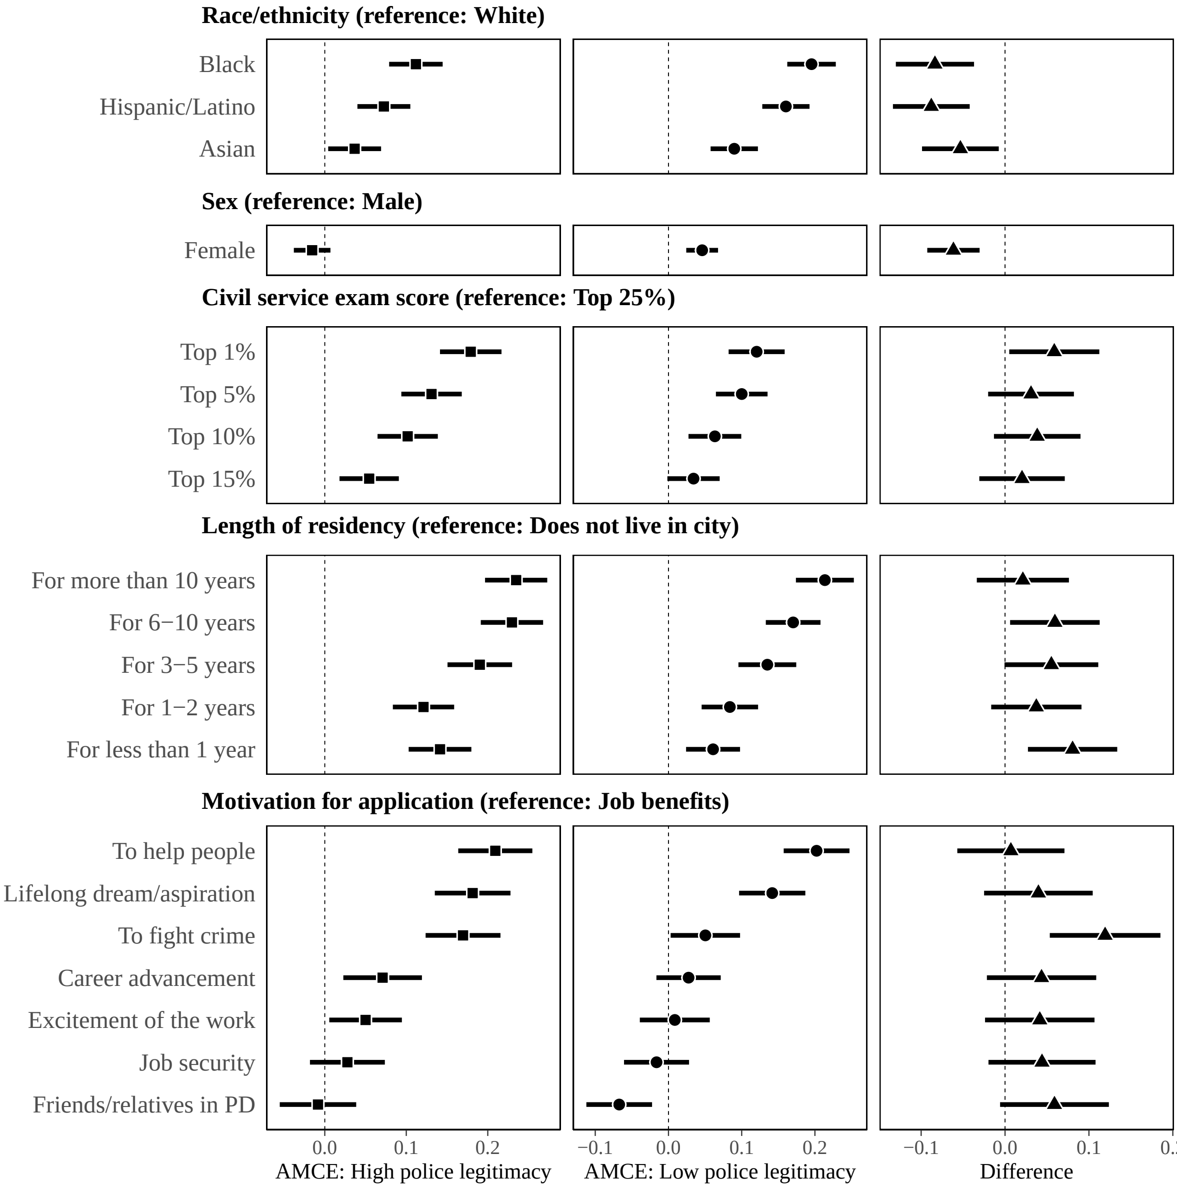

Figure S55: Estimated AMCEs in police recruitment conjoint by police legitimacy. Sub-group classifications are based on scores derived from the 10-item index of police legitimacy, trust, and confidence. Police legitimacy is coded as high if a respondent scored higher than the median respondent on the index. Point estimates and 95% confidence intervals estimated via OLS regression with robust standard errors clustered at respondent level to correct for within-respondent clustering. Municipal survey of Yonkers residents fielded in May 2021 (N = 1,413 respondents x 5 pairings x 2 applicants per pair = 14,130 observations).

S3 Pre-registration for information provision experiment

ASPREDICTED

Wharton CREDIBILITY LAB
UNIVERSITY of PENNSYLVANIA

CONFIDENTIAL - FOR PEER-REVIEW ONLY

Police diversity experiment on resident population, October 2021 (#76977)

Created: 10/14/2021 10:40 AM (PT)

This is an anonymized copy (without author names) of the pre-registration. It was created by the author(s) to use during peer-review.
A non-anonymized version (containing author names) should be made available by the authors when the work it supports is made public.

1) Have any data been collected for this study already?

No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?

Does information about demographic disparities (race/ethnicity and gender) between police and the residents they serve affect public support for policy change, trust, and willingness to cooperation with the police?

3) Describe the key dependent variable(s) specifying how they will be measured.

Behavioral outcomes: [1] Willingness to advocate for diversity policy (binary) by sending a message to local representative; [2] charitable contribution to pro-diversity organization from lottery payment (continuous $0-50). Attitude indices: [1] Support for diversification. Support for hiring underrepresented applicants (4 items): each item is a choice between two equally qualified candidates w/ three options: hire the underrepresented applicant (e.g., White), hire the other candidate (e.g., Black), or random selection (e.g., let a coin flip decide). Support for affirmative action (6 items): one item eliciting general support for implementing affirmative action programs at police department, and four items eliciting support for hiring from each underrepresented group, all on 7-point scale from Strongly Oppose to Strongly Support. One item that elicits rank ordering of affirmative action programs against alternatives (Body worn cameras, Civilian Review Boards, Community Policing Programs) with rank of affirmative action recorded on 4-point scale from least preferred to most preferred. [2] Trust in police: 1) How much of the time do you think [CITY NAME] residents can trust the [CITY NAME] Police Department to do what is right? (1-5 scale, Never-Always); 2) How much confidence do you have [CITY NAME] police to act in the best interest of the public? (1-5 scale, None-A great deal). [3] Willingness to cooperate: 1) How likely would you be to attend a community meeting to discuss problems in your neighborhood with the police?; 2) How likely would you be to report suspicious activity to the police?; 3) How likely would you be to call the police to report a crime?; 4) If the police were looking for a suspect who was hiding, and you knew where that person was, how likely would you be to provide the police with information? All items measured on 7-point scale, Extremely Unlikely to Extremely Likely. [4] Willingness to associate: 2-item scale eliciting willingness to 1) consider a career at [CITY NAME] police department; 2) encourage a friend/family member to consider a career at [CITY NAME] police department. Items measured on 7-point scale, Extremely Unlikely to Extremely Likely.

4) How many and which conditions will participants be assigned to?

2 groups. Treated respondents will be exposed to accurate information about demographic disparities between police department and community. Control respondents will receive no information about demographic disparities. All respondents will be asked to provide their best guess about the demographic representation of each group within the police department prior to randomization.

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.

1) Average treatment effects (ATEs) estimated via regression, with covariate adjustment to increase precision. Pre-treatment covariates include demographics (e.g., partisanship, race/ethnicity), respondents' over/under estimation of demographic disparities, as well as baseline attitudes toward police (e.g., trust/cooperation from baseline survey) and support for diversification from prior survey wave. 2) Conditional average treatment effects (CATEs) estimated among sub-groups defined by race/ethnicity (White v. non-White), sex (male v. female), partisanship, and whether respondents are over/under estimators of demographic disparities.

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.

Not applicable.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.

A baseline survey has been completed (N = 1,413), which contains respondent demographics and baseline measures of the outcomes referenced in 5). We anticipate ~1,000 respondents from the baseline survey will complete this survey. With N = 1,000, minimum detectable effect (.80 power) for attitudinal indices is ~0.17 standard units under conservative assumption of no precision gains from covariate-adjustment.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)

We will collect data on respondents' beliefs about diversity among U.S. police in general, their causal attributions for demographic disparities between police and communities, and perceived importance of minority representation among police. We will also examine treatment effect heterogeneity as a function of pre-treatment covariates via causal forests. All these analyses will be exploratory.

S4 Pre-registration for conjoint experiment

AS PREDICTED

Wharton CREDIBILITY LAB
UNIVERSITY of PENNSYLVANIA

CONFIDENTIAL - FOR PEER-REVIEW ONLY

YPD recruitment conjoint (#65678)

Created: 05/11/2021 11:59 AM (PT)

This is an anonymized copy (without author names) of the pre-registration. It was created by the author(s) to use during peer-review.
A non-anonymized version (containing author names) should be made available by the authors when the work it supports is made public.

---

1) Have any data been collected for this study already?
No, no data have been collected for this study yet.

2) What's the main question being asked or hypothesis being tested in this study?
This study uses a conjoint experiment to quantify how the attributes enumerated below affect preferences for hiring police officers. This is primarily a descriptive experiment that seeks to examine which attributes of police recruits are most influential in a multidimensional context.

3) Describe the key dependent variable(s) specifying how they will be measured.
Respondents will choose between two hypothetical applicants to the police department. The primary outcome is a binary choice between the persons presented during each conjoint task: "If you had to choose between them, which of these two applicants would you prefer to see recruited?". The secondary outcome is a 7-point scale: "Please rate each applicant on a scale from 1 to 7, where 1 indicates they should definitely not be recruited and 7 indicates they should definitely should be recruited."

4) How many and which conditions will participants be assigned to?
Each respondent will evaluate five conjoint tasks with eight attributes. The attributes (and levels): 1) Race/ethnicity ("White","Black","Hispanic/Latino","Asian"); 2) Sex ("Male","Female"); 3) Age ("23","25","27","29","31","33","35","37"); 4) Residency ("Does not live in City","For less than 1 year","For 1-2 years","For 3-5 years","For 6-10 years","For more than 10 years"); 5) Education ("GED","High school","Associates degree","Bachelors degree","Graduate degree"); 6) Civil service exam ("Scored in top 1% of applicants","Scored in top 5% of applicants","Scored in top 10% of applicants","Scored in top 15% of applicants","Scored in top 25% of applicants"); 7) Previous occupation ("Police officer in another city","Security guard","School teacher","Construction worker","Military service","Server/Bartender","Retail salesperson","Personal trainer","Social worker"); 8) Motivation for becoming a police officer ("Job benefits (i.e. medical/pension)","Excitement of the work","Opportunity to help people","To fight crime","Job security","Lifelong dream/aspiration","Has friends/relatives in police department","Opportunities for career advancement").

5) Specify exactly which analyses you will conduct to examine the main question/hypothesis.
The estimands of primary interest in this conjoint are the AMCEs. We will use linear regression of the outcome measures on the randomized attributes, with robust standard errors clustered at the respondent level. We will also report the marginal means for the levels within each attribute

6) Describe exactly how outliers will be defined and handled, and your precise rule(s) for excluding observations.
Not applicable.

7) How many observations will be collected or what will determine sample size? No need to justify decision, but be precise about exactly how the number will be determined.
Approximately 2500 participants. The experiment will be administered to about 2000 residents of Yonkers, NY as part of a community survey. The experiment will also be administered to approximately 500 officers in the Yonkers Police Department. These counts are based on estimated response rates to the surveys.

8) Anything else you would like to pre-register? (e.g., secondary analyses, variables collected for exploratory purposes, unusual analyses planned?)
We will conduct exploratory sub-group analyses based on respondents' partisanship, race/ethnicity, gender, and attitudes toward police (trust/legitimacy).

Version of AsPredicted Questions: 2.00

Available at https://aspredicted.org/WGM_QGW

References

[1] Coppock A, McClellan OA (2019) Validating the demographic, political, psychological, and experimental results obtained from a new source of online survey respondents. *Research & Politics* 6(1):2053168018822174.

[2] Berinsky AJ, Huber GA, Lenz GS (2012) Evaluating online labor markets for experimental research: Amazon. com's mechanical turk. *Political analysis* 20(3):351–368.

[3] Mullinix KJ, Leeper TJ, Druckman JN, Freese J (2015) The generalizability of survey experiments. *Journal of Experimental Political Science* 2(2):109–138.

[4] Coppock A, Leeper TJ, Mullinix KJ (2018) Generalizability of heterogeneous treatment effect estimates across samples. *Proceedings of the National Academy of Sciences* 115(49):12441–12446.

[5] Coppock A (2019) Generalizing from survey experiments conducted on mechanical turk: A replication approach. *Political Science Research and Methods* 7(3):613–628.

[6] Peyton K, Huber GA, Coppock A (2021) The generalizability of online experiments conducted during the covid-19 pandemic. *Journal of Experimental Political Science* p. 1–16.

[7] Berinsky AJ, Margolis MF, Sances MW (2014) Separating the shirkers from the workers? making sure respondents pay attention on self-administered surveys. *American Journal of Political Science* 58(3):739–753.

[8] Aronow PM, Kalla J, Orr L, Ternovski J (2020) Evidence of rising rates of inattentiveness on lucid in 2020.

[9] Ternovski J, Orr L (2022) A note on increases in inattentive online survey-takers since 2020. *Journal of Quantitative Description: Digital Media* 2.

[10] Montgomery JM, Nyhan B, Torres M (2018) How conditioning on posttreatment variables can ruin your experiment and what to do about it. *American Journal of Political Science* 62(3):760–775.

[11] Aronow PM, Baron J, Pinson L (2019) A note on dropping experimental subjects who fail a manipulation check. *Political Analysis* 27(4):572–589.

[12] Vaughn PE, Peyton K, A. HG (2022) Mass support for proposals to reshape policing depends on the implications for crime and safety. *Criminology & Public Policy* 21(1):125–146.

[13] Rudkin A (2022) *autumn: Fast, Modern, and Tidy-Friendly Iterative Raking.* R package version 0.1.

[14] DeBell M, Krosnick JA (2009) Computing weights for american national election study survey data. *nes012427. Ann Arbor, MI, Palo Alto, CA: ANES Technical Report Series.*

[15] Miratrix LW, Sekhon JS, Theodoridis AG, Campos LF (2018) Worth weighting? how to think about and use weights in survey experiments. *Political Analysis* 26(3):275–291.

[16] Broockman D, Kalla J (2016) Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science* 352(6282):220–224.

[17] Broockman DE, Kalla JL, Sekhon JS (2017) The design of field experiments with survey outcomes: A framework for selecting more efficient, robust, and ethical designs. *Political Analysis* 25(4):435–464.

[18] Yan A, Kalla J, Broockman DE (2018) Increasing response rates and representativeness of online panels recruited by mail: Evidence from experiments in 12 original surveys.

[19] Peyton K, Sierra-Arevalo M, Rand DG (2019) A field experiment on community policing and police legitimacy. *Proceedings of the National Academy of Sciences* 116(40):19894–19898.

[20] Nix J, Pickett JT, Baek H, Alpert GP (2019) Police research, officer surveys, and response rates. *Policing and Society* 29(5):530–550.

[21] Ba BA, Knox D, Mummolo J, Rivera R (2021) The role of officer race and gender in police-civilian interactions in chicago. *Science* 371(6530):696–702.

[22] Anderson ML (2008) Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association* 103(484):1481–1495.

[23] Raganella AJ, White MD (2004) Race, gender, and motivation for becoming a police officer: Implications for building a representative police department. *Journal of criminal justice* 32(6):501–513.

[24] White MD, Cooper JA, Saunders J, Raganella AJ (2010) Motivations for becoming a police officer: Re-assessing officer attitudes and job satisfaction after six years on the street. *Journal of criminal justice* 38(4):520–530.

[25] Gibbs JC (2019) Diversifying the police applicant pool: Motivations of women and minority candidates seeking police employment. *Criminal Justice Studies* 32(3):207–221.

[26] Hainmueller J, Hopkins DJ, Yamamoto T (2014) Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis* 22(01):1–30.

[27] Hainmueller J, Hopkins DJ (2015) The hidden american immigration consensus: A conjoint analysis of attitudes toward immigrants. *American Journal of Political Science* 59(3):529–548.

[28] Glass GV (1976) Primary, secondary, and meta-analysis of research. *Educational researcher* 5(10):3–8.

[29] Gerber AS, Green DP (2012) *Field experiments: Design, analysis, and interpretation.* (WW Norton).

[30] Lin W (2013) Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *The Annals of Applied Statistics* 7(1):295–318.

[31] Wager S, Du W, Taylor J, Tibshirani RJ (2016) High-dimensional regression adjustments in randomized experiments. *Proceedings of the National Academy of Sciences* 113(45):12673–12678.

[32] Center PR (2019) Why americans don't fully trust many who hold positions of power and responsibility.

[33] Wager S, Athey S (2018) Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113(523):1228–1242.

[34] Athey S, Wager S (2019) Estimating treatment effects with causal forests: An application. *Observational Studies* 5(2):37–51.

[35] Athey S, Tibshirani J, Wager S (2019) Generalized random forests. *The Annals of Statistics* 47(2):1148–1178.

[36] Tibshirani J, Athey S, Sverdrup E, Wager S (2021) *grf: Generalized Random Forests.* R package version 2.0.2.

[37] Guess A, Coppock A (2018) Does counter-attitudinal information cause backlash? *British Journal of Political Science* pp. 1–19.

[38] Peyton K (2020) Does trust in government increase support for redistribution? evidence from randomized survey experiments. *American Political Science Review* 114(2):596–602.

[39] Hernán MA, Robins JM (2006) Instruments for causal inference. *Epidemiology* pp. 360–372.

[40] Sovey AJ, Green DP (2011) Instrumental variables estimation in political science. *American Journal of Political Science* 55(1):188–200.

[41] Stock JH, Wright JH, Yogo M (2002) A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics* 20(4):518–529.

[42] Angrist JD, Pischke JS (2009) *Mostly Harmless Econometrics: An Empiricist's Companion.* (Princeton: Princeton University Press).