

**Beliefs about racial discrimination and support
for pro-black policies**

Ingar Haaland & Christopher Roth

[\(This paper also appears as CAGE Discussion paper 554\)](#)

March 2021

No: 1339

Warwick Economics Research Papers

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

Beliefs About Racial Discrimination and Support for Pro-Black Policies

Ingar Haaland Christopher Roth*

March 31, 2021

Abstract—This paper provides representative evidence on beliefs about racial discrimination and examines whether information causally affects support for pro-black policies. Eliciting quantitative beliefs about the extent of hiring discrimination against blacks, we uncover large disagreement about the extent of racial discrimination with particularly pronounced partisan differences. An information treatment leads to a convergence in beliefs about racial discrimination but does not lead to a similar convergence in support of pro-black policies. The results demonstrate that while providing information can substantially reduce disagreement about the extent of racial discrimination, it is not sufficient to reduce disagreement about pro-black policies.

Keywords: Racial Discrimination, Beliefs, Pro-Black Policies, Policy Preferences

JEL Codes: C91, D83, J71, J15

*Haaland: University of Bergen and CESifo; Roth: University of Warwick, briq, CESifo, CEPR, and CAGE. We would like to thank the editor (Ray Fisman) and four anonymous reviewers for excellent and highly constructive suggestions. We also would like to thank Roland Bénabou, Björn Bartling, Alexander W. Capellen, Eliana La Ferrara, Paola Giuliano, Matt Lowe, Gautam Rao, Bertil Tungodden and seminar participants in Bergen, Bonn, Cologne (EEA), Copenhagen, Oslo, Oxford, Mannheim, and Munich for helpful comments and discussions. We thank Hrishikesh Iyengar for excellent research assistance. Financial support from the Russell Sage Foundation (Small Awards in Behavioral Economics), the Research Council of Norway through its Centre of Excellence Scheme (FAIR project No 262675), and data collection by Time-sharing Experiments for the Social Sciences (NSF Grant 0818839, Jeremy Freese and James Druckman, Principal Investigators) is gratefully acknowledged. All six experiments in this study were pre-registered (AEA RCT Registry trial 2273). IRB approvals were obtained from the University of Oxford and the NHH Norwegian School of Economics. The usual disclaimer applies.

I. Introduction

Racial discrimination is a pervasive phenomenon that affects many spheres of society (Arrow, 1998; Bertrand and Duflo, 2017; List, 2004). In the United States, several studies have documented high levels of racial discrimination in various domains, such as the labor market (Bertrand and Mullainathan, 2004), the housing market (Edelman et al., 2017), and the judicial system (Alesina and La Ferrara, 2014).

To deal with this large degree of racial discrimination, the US government has introduced policies aiming to actively counteract the effects of racial discrimination. However, Americans are deeply divided in their support for such policies. For instance, while 73 percent of Democrats support affirmative action programs for racial minorities, only 38 percent of Republicans support this.¹ There is a strong perception in the public debate that this political disagreement is rooted in differences in perceptions of the extent of racial discrimination in society (Newkirk, 2017). Similarly, in a seminal article on the drivers of opposition to pro-black policies, Bobo and Kluegel (1993) argue that it is necessary to correct people’s biases in beliefs to gain support for pro-black policies.

In this paper, we examine the relationship between beliefs about racial discrimination and support for pro-black policies. We address the following three main questions: First, how large is the disagreement in beliefs about the extent of racial discrimination in hiring? Second, does the provision of research evidence about the extent of racial discrimination in hiring reduce the disagreement about the extent of racial discrimination? Third, does a convergence in beliefs about the extent of racial discrimination lead to a similar convergence in support for pro-black policies?

To study these questions, we introduce a new approach to elicit quantitative beliefs about racial discrimination. With respondents from a high-quality, probability-based sample of the US household population, we elicited incentivized beliefs about the results of a correspondence study testing for racial discrimination against blacks in the labor market (Bertrand and Mullainathan, 2004). Respondents were told that researchers sent out

¹ <https://news.gallup.com/poll/184772/higher-support-gender-affirmative-action-race.aspx> (accessed September 29, 2020).

resumes that were identical in all respects except for the perceived race of the sender to help wanted ads in Boston and Chicago newspapers. After informing the respondents that resumes with white-sounding names had to be sent out ten times to get one callback on average, we asked them how many times they thought that resumes with black-sounding names had to be sent out to get one callback on average.

To examine whether information about racial discrimination in hiring causally affects people's support for policies aiming to counteract the effects of racial discrimination in hiring, we introduced exogenous variation in people's beliefs by informing a random subset of the respondents about the actual results from the correspondence study by Bertrand and Mullainathan (2004), namely that white-sounding names received 50 percent more callbacks for interviews than black-sounding names. To measure whether people update their beliefs about racial discrimination in response to this evidence, we elicited their beliefs about a second correspondence study that tested for racial discrimination in the housing market (Edelman et al., 2017). Furthermore, to measure whether the information provision affects people's political behavior, respondents decided whether to receive money versus making a real donation to a pro-black civil rights organization. Finally, respondents answered a series of questions on self-reported views on pro-black policies.

We document several novel findings on beliefs about racial discrimination and support for pro-black policies in the United States. Our first finding is that there is substantial disagreement about the extent of racial discrimination against blacks. As illustrated in Figure A.1, we find particularly pronounced differences in beliefs based on people's self-identified party affiliation: Relative to Republicans, Democrats think that blacks have to send out 47 percent more resumes than whites to receive a callback. Second, our respondents strongly update their beliefs in response to research evidence from a correspondence study, substantially reducing disagreement about the extent of racial discrimination. Third, we find a muted response of the information on people's support for pro-black policies, demonstrating that a strong convergence in beliefs about racial discrimination does not necessarily lead to a convergence in policy views. Finally, using an additional set of experiments, we examine some potential mechanisms behind our results

and demonstrate the robustness of our findings to designs which mitigate concerns about social desirability bias and question framing.

We have two main contributions. First, we introduce a new approach for measuring incentivized and quantitative beliefs about racial discrimination by eliciting predictions about the results from correspondence studies. In contrast to traditional survey questions, which typically ask about “how much discrimination is there” on a scale from “a lot” to “none at all,” our approach allows us to elicit quantitative and incentivized beliefs about racial discrimination in a precisely defined environment. Our paper thus relates to the literature on political polarization of beliefs (Alesina et al., 2020, 2018b; Allcott et al., 2020; Grigorieff et al., 2020; Kuziemko et al., 2015). We contribute to this literature by providing the first quantitative and incentivized data which allows for measuring political belief polarization in the context of racial discrimination. Our approach is also related to an emerging literature on the measurement of prior beliefs about the results of research studies (DellaVigna and Pope, 2018a,b; DellaVigna et al., 2019).

Second, we provide the first causal evidence on how beliefs about racial discrimination affect the demand for policies that try to counteract the effects of discrimination. Thematically, our paper relates closely to work in political science on the determinants of support for pro-black policies (Bobo and Kluegel, 1993; Harrison et al., 2006; Kluegel and Smith, 1983; Kuklinski et al., 1997). Our primary finding is that a convergence in beliefs about racial discrimination does not result in a convergence in support for pro-black policies. We thereby relate to a literature studying how information provision affects people’s policy preferences (see, e.g., Cruces et al., 2013; Karadja et al., 2017; Kuziemko et al., 2015) and to a recent literature on how people update their beliefs from research findings (Haaland and Roth, 2020; Hjort et al., 2020; Vivald and Coville, 2020).² In a closely related study, Haaland and Roth (2020) study how the provision of research evidence about the labor market impact of immigration affects beliefs and attitudes towards immigration. They find that a convergence in beliefs about the labor market impact of immigration leads to

² See also Haaland et al. (2020) for a recent literature review of information provision experiments with additional citations.

a similar convergence in attitudes towards immigration. Taken together with the findings from the present study, where a convergence in beliefs does not lead to a convergence in support for pro-black policies, the results suggest that the relationship between beliefs and policy preferences is context specific. Finally, our findings complement previous work examining whether the awareness of discrimination reduces biased judgments (Alesina et al., 2018a; Pope et al., 2018).

The remainder of the paper proceeds as follows. Section II describes the main experimental design and the main sample. Section III provides descriptive data on people’s beliefs about racial discrimination and presents treatment effects of the provision of research evidence about the extent of racial discrimination against blacks on beliefs and policy views. Section IV provides additional evidence on robustness and mechanisms. Section V concludes. The Online Appendix provides additional results and the full set of experimental instructions.

II. Experimental design: Main Experiment

In this section, we describe the sample and experimental design of the main experiment (Experiment 1).

A. Sample

We recruited 1382 respondents through NORC’s AmeriSpeak panel in June and July 2017.³ AmeriSpeak is a probability-based panel of the US population. The panel uses NORC’s National Frame, which is designed to provide at least 97 percent sample coverage of the US population. The NORC National Frame is used for several landmark studies in

³ 1542 respondents completed any of the outcome questions, but only 1382 of these respondents gave an answer to the pre-treatment question on the number of times resumes with black-sounding names had to be sent (NORC does not force their respondents to answer any questions on their surveys). In the main analysis, we therefore focus on the 1382 respondents who answered this question. There are no significant differences between Republicans and Democrats or between blacks and whites in not responding to this question. Our main specification includes only respondents who completed the question on beliefs about racial discrimination.

the US, including the General Social Survey (GSS), which is one of the most frequently-analyzed data sets in the social sciences. Table A.2 provides summary statistics for this sample. Observations in the treatment and control group are balanced in terms of observables (Table A.3).

B. Design

Structure The structure of Experiment 1 is as follows (Figure A.2 provides an overview): We first measured beliefs about racial labor market discrimination and then exposed half of our respondents to an information treatment. Subsequently, we measured people’s support for policies to address racial discrimination and also elicited post-treatment beliefs about racial discrimination.

Pre-treatment beliefs about racial labor market discrimination We used a correspondence study to measure people’s beliefs about racial discrimination in the labor market. Correspondence studies rely on fictitious resumes to study discrimination in the labor market (Bertrand and Duflo, 2017). Specifically, by manipulating whether a fictitious resume is assigned a minority name, researchers can study racial labor market discrimination by comparing the outcomes for resumes with and without the perceived minority name. A seminal correspondence study by Bertrand and Mullainathan (2004) found that white-sounding names were 50 percent more likely to receive a callback than black-sounding names; a finding that has been closely replicated in several subsequent correspondence studies (Bertrand and Duflo, 2017; Quillian et al., 2017). We rely on this study in our experiment. To familiarize our respondents with the study, we presented them with the following text:

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did so by sending out fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names

like “Tanisha” and “Kareem”. The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer whether they were white or black.

We informed respondents that resumes with white-sounding names had to be sent out on average ten times to get one callback for an interview. To measure their beliefs about racial discrimination in the labor market, we then asked how many times they believe resumes with black-sounding names had to be sent out on average to get one callback for an interview. Furthermore, we promised respondents a \$2 bonus if their answer was the same “as what the researchers found.”

Our belief elicitation has several advantages compared to qualitative survey questions that have traditionally been used to study beliefs about racial discrimination. First, we measure beliefs on a quantitative scale that is easily comparable across respondents and has the same interpretation for everyone. By contrast, many previous studies have assessed beliefs about racial discrimination using a question from the General Social Survey about the extent of discrimination that blacks face in “getting good jobs,” which is measured on a 4-point scale from “none at all” to “a lot.”⁴ One concern with using subjective response scales to measure beliefs is that different people may have different opinions about what, e.g., “some” or “only a little” discrimination means.⁵ Furthermore, in our setting, racial discrimination is precisely defined, allowing us to hold our respondents’ beliefs about the circumstances of racial discrimination constant. For qualitative survey questions, people may hold different beliefs about what constitutes “discrimination.” These beliefs might be correlated with demographics, which makes it difficult to draw strong conclusions on differences in beliefs about racial discrimination across demographic groups. Our measure avoids these confounds. Second, non-incentivized survey questions in the political domain are potentially more prone to the misreporting of beliefs. Indeed, small incentives for

⁴ Details about this variable are available at the following link: <https://gssdataexplorer.norc.org/variables/1244/vshow> (accessed November 30, 2018).

⁵ For a discussion of problems associated with subjective response scales, see Bond and Lang (2019).

correct answers have been shown to strongly increase the accuracy of survey responses and to reduce gaps in reported beliefs across party lines (Bullock et al., 2015; Prior et al., 2015). Since our question has a factual answer, we can incentivize correct responses.

Introducing exogenous variation in beliefs Two central identification challenges when studying the impact of beliefs on policy preferences are omitted variable bias and reverse causality. We address these identification challenges by introducing exogenous variation in beliefs, namely by informing respondents in the treatment group about the extent of racial discrimination found in the study by Bertrand and Mullainathan (2004). Specifically, we showed the following text to treated respondents:

The researchers found that resumes with black-sounding names on average had to be sent out 15 times to get one callback for an interview.

Since resumes with white-sounding names on average only had to be sent out 10 times to get one callback for an interview, this means that employers were 50 percent more likely to give callbacks to applicants with white-sounding names compared to applicants with black-sounding names.

By contrast, respondents in the control group did not receive any information and proceeded directly from the belief elicitation to the outcome questions.

Measuring support for pro-black policies: Behavioral measure A common critique of self-reported survey questions is that they might not be reflective of real political behavior and that they are prone to experimenter demand effects. To address these concerns, we collected a behavioral outcome measure, namely real donations to a pro-black civil rights organization, *The Lawyers' Committee for Civil Rights*. We told our respondents that they have the opportunity to financially support a civil rights organization that works to reduce discrimination against blacks in the labor market. We elicited the respondents' marginal rate of substitution between money for themselves and money for the civil rights organization through a multiple price list. The respondents chose between donating \$5 to the civil rights organization and money for themselves in \$1-increments from \$0 to \$5. One of the six choices was randomly implemented.

Measuring support for pro-black policies: self-reported policy views In addition to the behavioral measure, we also collected some data on people’s self-reported policy views. Since our treatment was tailored to shift beliefs about racial discrimination in the labor market, we focused on labor market policies. We asked questions about three commonly-discussed policies attempting to counteract the effects of labor market discrimination. First, we asked respondents whether they “support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job.” Second, we asked respondents whether they “support or oppose government and private programs that give qualified black candidates assistance in getting a job.” Third, we asked respondents whether they “support or oppose mandatory name-blind recruitment for hiring in public and private jobs.” For all three questions, respondents reported their answer on a 5-point scale ranging from 1 (Strongly oppose) to 5 (Strongly support).

Measuring beliefs about racial discrimination in the housing market To measure whether respondents updated their beliefs in response to the research evidence, we relied on a second correspondence study that tested for racial discrimination in the housing market (Edelman et al., 2017). We chose to focus on racial discrimination in a different domain out of a concern that demand effects, numerical anchoring, or a taste for consistency in survey responses could bias responses if we re-asked the question about discrimination in the labor market shortly after the information provision. The housing market is a good candidate for several reasons. First, racial discrimination in the housing market holds strong economic importance. Second, the study by Edelman et al. (2017), which serves as our benchmark for incentivizing beliefs, used the same names and general methodology as Bertrand and Mullainathan (2004). We used the following text to familiarize our respondents with the second study:

Researchers from Harvard Business School conducted an experiment to study racial discrimination in the rental market by sending out reservation requests from invented accounts to hosts on Airbnb, a website for private rental accommodations. The requests were exactly the same except for one thing: the name of the person

who sent the request. Half of the requests came from typically white-sounding names, while the other half came from typically black-sounding names. The idea was that the hosts would use the applicants' name to infer whether the reservation requests came from white or black requesters.

We then told them that the researchers found that white-sounding names were accepted 49 percent of the time. To measure their beliefs about racial discrimination in the housing market, we then asked what percent of the time they believe that black-sounding names were accepted. We offered a \$2 bonus for answers that fall within “2 percentage points of what the researchers found.”

We purposefully designed the second belief elicitation to avoid potential bias stemming from numerical anchoring by (i) using a different response scale than the first belief elicitation, and (ii) using a scale in which higher values implied less racial discrimination. Since higher values implied more discrimination in the first belief elicitation, numerical anchoring would make finding evidence for belief updating in the expected direction less likely.

III. Results: Main Experiment

This section uses data from Experiment 1 with NORC to provide representative evidence of people's beliefs about racial discrimination and to study the effects of the information treatment on support for pro-black policies.

A. *Heterogeneity in beliefs about racial discrimination*

Distribution of beliefs Figure 1 provides representative evidence of people's beliefs about racial discrimination in the labor and housing markets. Panel A shows the cumulative distribution function for beliefs about how many resumes with black-sounding names had to be sent out to get one callback on average (respondents were told that the corresponding number for white-sounding names was ten). This quantitative belief elicitation allows us to assess the fraction of respondents who overestimate and underestimate racial discrimination using the results from Bertrand and Mullainathan (2004) as the benchmark.⁶ Relative to

⁶ The study by Bertrand and Mullainathan (2004) is the most natural benchmark to use since our respondents were made familiar with the context of that specific study. As is

this study—which found that resumes with black-sounding names needed to be sent out 15 times before receiving one callback on average—we find that 35 percent of our respondents underestimate racial discrimination in the labor market, 10.3 percent have correct beliefs, and the remaining 54.7 percent overestimate the extent of racial discrimination in the labor market.

Panel B of Figure 1 shows the cumulative distribution function for beliefs about the rejection rate of reservation requests from black-sounding names on Airbnb (respondents were told that the corresponding number for white-sounding names was 51 percent). Using the results from Edelman et al. (2017)—who found that requests from black-sounding names were rejected 59 percent of the time—as the benchmark, we find that 19 percent of our respondents underestimate racial discrimination in the housing market and the remaining 81 percent overestimate the extent of racial discrimination in the housing market.

The data also allows for the measurement of the share of respondents who think that there is discrimination against whites, discrimination against blacks, and no racial discrimination at all. For the labor market, 23 percent of our respondents believe that there is discrimination against whites, nine percent believe that there is no discrimination, and the remaining 68 percent believe that there is discrimination against blacks. For the housing market, 12 percent think that there is discrimination against whites, two percent believe that there is no racial discrimination, and the remaining 86 percent think that there is discrimination against blacks.

[Insert Figure 1 here]

true for all results obtained from a single study, this benchmark is a noisy estimate of the truth which depends on sampling variation as well as contextual factors. While a recent meta-analysis provides evidence of no change in racial discrimination in hiring over time (Quillian et al., 2017), there is substantial variation in estimates of racial discrimination when comparing individual studies. The analysis in this section on the fraction of Americans who overestimate or underestimate racial discrimination is therefore suggestive and should not be interpreted as conclusive evidence of misperceptions.

Correlates of beliefs Figure 2 examines whether beliefs about racial discrimination vary systematically by people’s background characteristics. Panel A shows correlations between background characteristics and beliefs about racial discrimination in the labor market. We find especially pronounced differences in beliefs based on people’s political affiliation: Relative to Republicans, Democrats believe that seven additional resumes with black-sounding names had to be sent out to get one callback on average ($p < 0.01$).⁷ Relative to those with no college education, college-educated respondents believe that four additional resumes with black-sounding names had to be sent out to get one callback on average ($p < 0.01$). Relative to respondents with below median income, above-median income respondents believe that 1.7 additional resumes with black-sounding names had to be sent out to get one callback on average ($p < 0.05$). Surprisingly, we find no significant differences between blacks and whites in their beliefs about discrimination in the labor market ($p = 0.85$).

[Insert Figure 2 here]

Concerning beliefs about the housing market (Panel B of Figure 2), we also find pronounced differences based on people’s political affiliation: Relative to Republicans, Democrats think that reservation requests from black-sounding names were 5.7 percentage points more likely to be rejected ($p < 0.01$). While we do not find evidence of differences in beliefs in the housing market across people with different education levels, we find significant racial differences: Relative to whites, blacks think that reservation requests from black-sounding names were 6.5 percentage points more likely to be rejected ($p < 0.05$).

The association between beliefs and policy preferences Table 1 provides evidence using control group respondents in Experiment 1 on whether our measure of beliefs about racial labor discrimination correlates with some of our key outcome measures. Column 1 of Panel

⁷ Differences in beliefs between Republicans and Democrats might be driven by differences in information sets resulting from the coverage of different studies by right-leaning versus left-leaning news outlets. For example, Republicans may be more likely to hear about studies which find little evidence of racial discrimination, while Democrats may be more likely to hear about studies which find strong evidence on racial discrimination.

A shows a regression of people’s actual donations to the pro-black civil rights organization on their beliefs about racial discrimination in the labor market. A one standard deviation increase in beliefs is associated with 0.22 of a standard deviation higher donations to the pro-black civil rights organization ($p < 0.01$). This corresponds to 36 percent of the Democrat–Republican difference in donations to the pro-black civil rights organization. Including controls in the regression reduces the estimated association to 0.17 of a standard deviation ($p < 0.01$, Column 1 of Panel B). Columns 2 and 3 of Table 1 show significant associations between beliefs about racial discrimination and support for preference in hiring and job assistance for blacks, respectively. Column 4 shows that a one standard deviation change in beliefs about racial discrimination in the labor market is associated with a 0.22 of a standard deviation change in beliefs about discrimination in the housing market. Furthermore, column 5 shows that our belief measure is also predictive of whether people think that racial discrimination against blacks in the labor market is a “serious problem.”

Our first main result based on the findings in this section can be summarized as follows:

Result 1. *There is a large disagreement about the extent of racial discrimination with particularly pronounced partisan differences. Beliefs about the extent of racial discrimination are strongly associated with donations to a pro-black civil rights organization and self-reported support for pro-black policies.*

[Insert Table 1 here]

B. Treatment effects on beliefs and policy views

Empirical strategy We pre-specified the analysis in a document uploaded to the AEA RCT Registry prior to starting the data collection. The empirical strategy outlined in this section follows the pre-analysis plans, which may be accessed with the following link: <https://www.socialscienceregistry.org/trials/2273>.

Since we expect different treatment effects based on whether the respondents initially overestimate or underestimate racial discrimination, our main specification is the following

equation, which we estimate using OLS:

$$y_i = \alpha_0 + \alpha_1 \text{Treatment}_i + \alpha_2 \text{Treatment}_i \times \text{overestimate}_i + \alpha_3 \text{overestimate}_i + \alpha_4 \mathbf{x}_i + \varepsilon_i$$

where y_i is the outcome of interest; Treatment_i is an indicator for whether respondent i received the research evidence; overestimate_i is an indicator for initially overestimating discrimination (i.e., for having pre-treatment beliefs that resumes with black-sounding names had to be sent out more than 15 times to get one callback on average)⁸; \mathbf{x}_i is a vector of pre-specified controls (gender dummy; age in years; two ethnicity indicators (non-Hispanic whites and non-Hispanic blacks); three region dummies; household size; log household income; an indicator for being employed; two party affiliation indicators (Republicans and Democrats); and an indicator for having a college degree); and ε_i is an individual-specific error term. We use robust standard errors for inference. Throughout the section, we refer to respondents who initially underestimate and overestimate racial discrimination in the labor market as “underestimators” and “overestimators,” respectively. *Do people update their beliefs about racial discrimination?* We first examine whether people use the information about racial discrimination in the labor market to update their beliefs about racial discrimination in the housing market.⁹ Figure 3 demonstrates that while there is a large and significant association between beliefs about discrimination in the labor market and in the housing market among control group respondents, there is no significant correlation among treated respondents. Figure 4 shows a similar convergence in beliefs about housing market discrimination when we compare overestimators and underestimators. More specifically, as shown in column 1 of Table 2, treated underestimators increase their

⁸ Since those with accurate pre-treatment beliefs (i.e., 15) should become more confident in their beliefs, and thus should increase support for pro-black policies, we decided to group them in the same category as those who strictly underestimated racial discrimination.

⁹ While respondents were asked about the acceptance rate of black-sounding names (i.e., what percent of the time they thought reservation requests from black-sounding names were accepted), we recoded the responses such that higher numbers imply more discrimination.

estimate of the rejection rate of black-sounding names by 4.2 percentage points (or 21 percent of a standard deviation; $p < 0.01$) while treated overestimators decrease their estimate of the rejection rate of black-sounding names by 5.8 percentage points (or 29 percent of a standard deviation; $p < 0.01$).

To assess the degree of learning from the information, we also examine heterogeneous treatment effects on posterior beliefs about discrimination in the housing market, posterior_i , by an individual’s perception_gap_i , defined as the difference in people’s prior beliefs and the signal they receive, i.e. $\text{prior}_i - 15$. Specifically, we estimate the following equation:

$$\text{posterior}_i = \beta_0 + \beta_1 \text{Treatment}_i + \beta_2 \text{Treatment}_i \times \text{perception_gap}_i + \beta_3 \text{perception_gap}_i + \beta_4 \mathbf{x}_i + \varepsilon_i$$

As shown in column 2 of Table 2, treated respondents increase their estimate of the rejection rate among blacks by 0.37 percentage points in response to learning that blacks had to send out one more CV than previously thought ($p < 0.01$).¹⁰ There is no significant treatment heterogeneity in the extent of belief updating between Republicans and non-Republicans, but looking at point estimates, we find suggestive evidence of higher updating among Republicans than among non-Republicans for a given perception gap (Columns 1 and 2 of Table A.8).¹¹ Furthermore, using the results from Edelman et al. (2017) as the

¹⁰ We find no significant differences in learning rates among overestimators and underestimators.

¹¹ As described in the pre-analysis plan, we made an ex ante decision to focus on heterogeneity between Republicans and non-Republicans. We did this for two main reasons. First, political affiliation is a stronger predictor of support for pro-black policies than all other demographics, including race (<http://pewrsr.ch/2wAjUGP>; accessed February 4, 2019). Second, the Republican/non-Republican split allows us to include all respondents in the heterogeneity analysis, maximizing statistical power. Third, there are theoretical reasons for why political affiliation may be an important moderator of treatment effects. On the one hand, the literature on motivated reasoning shows that people who receive information that goes against their political convictions might be less willing to update their beliefs than people for whom the information is in line with

benchmark, we find that treated respondents on average hold less biased beliefs about housing market discrimination compared to control group respondents (columns 1 to 3 of Table A.9). Given all of the findings discussed above, our second main result can be summarized as follows:

Result 2. *Americans strongly update their beliefs about racial discrimination in response to research evidence from correspondence studies. Provision of the research evidence leads to a convergence in beliefs about housing market discrimination between those who initially overestimate and underestimate racial discrimination in the labor market.*

[Insert Figure 3 here]

[Insert Table 2 here]

Does the information affect support for pro-black policies? Table 3 presents regressions results for our main outcomes of interest. All outcomes are z-scored using the mean and standard deviation in the control group. Column 2 shows treatment effects on donations to the pro-black civil rights organization. While treated underestimators increase their donations by 0.16 of a standard deviation ($p < 0.05$), the treatment effect on overestimators is close to zero and not statistically significant ($p = 0.79$). Furthermore, columns 3-6 show insignificant treatment effects close to zero on self-reported policy views for both overestimators and underestimators. Overall, we thus observe a muted impact of treatment on support for pro-black policies even though the treatment led to a strong convergence in beliefs between overestimators and underestimators (as shown in Figure 4).¹²

We next examine political heterogeneity in treatment effects. While we do not detect their political orientation (Taber and Lodge, 2006). On the other hand, there could be a ceiling effect given that Democrats in the United States are more supportive of pro-black policies to being with, leaving less room to change their policy preferences.

¹² There are some important differences between the donation measure and self-reported measures. First, the donation measure involves a real trade-off between money for the respondent versus money for civil rights organization, making it a more externally valid outcome measure than the self-reports. Second, the donation behavior depends on people's beliefs about the efficiency of organization in addressing racial discrimination,

any significant political heterogeneity in donation behavior, the patterns in the data are generally consistent with stronger treatment effects for non-Republicans and weaker treatment effects for Republicans (as shown in columns 3 to 6 of Table A.11). The muted response among Republican underestimators is especially striking considering that this group most strongly updated their beliefs in response to the information (column 1 of Table A.8), but heterogeneous treatment effects along party lines are not precisely estimated due to small cell sizes.¹³ On self-reported policy views, we generally see muted effects close to zero for all subgroups and outcomes (columns 1 to 4 of Table A.14). It is thus not the case that the overall muted impact of the information on policy preferences masks important political heterogeneity going in opposite directions.¹⁴ Our third main result can thus be summarized as follows:

Result 3. *Support for pro-black policies is largely unresponsive to information about the extent of discrimination against blacks in the labor market.*

[Insert Figure 4 here]

[Insert Table 3 here]

IV. Robustness Experiments and Potential Mechanisms

Experiment 1 demonstrated that a strong convergence in beliefs about racial discrimination did not lead to a similar convergence in support of pro-black policies. In this section, we present additional evidence to address potential concerns about social making heterogeneous treatment effects more challenging to interpret than for the self-reports.

¹³ One potential explanation for the heterogeneity could be differences in beliefs about the efficacy of the civil rights organization in addressing racial discrimination.

¹⁴ One explanation for the lack of response to the information among both Republican and non-Republican overestimators could be that they think that discrimination is still sufficiently prevalent to justify support for pro-black policies even though they were informed that discrimination is less prevalent than they thought. An additional potential explanation for the muted response among overestimators is that they already held qualitatively correct beliefs that there is discrimination against blacks.

desirability bias (Experiment 2) and robustness to framing and incentives (Experiment 3). We also examine beliefs about the sources of discrimination as an explanation for the muted treatment effects (Experiment 4) and whether support for pro-black policies is elastic to other types of information (Experiments 5 and 6). All additional experiments were pre-specified in the same AEA RCT Registry trial as the main experiment.

A. Experiment 2: Addressing social desirability bias

Motivation A potential concern with collecting the main outcome measures immediately after administering the treatment, as in Experiment 1, is that social desirability bias might differ between the treatment and control group (de Quidt et al., 2018). For instance, respondents may think that it is socially undesirable to express low support in pro-black policies immediately after they have been informed that whites receive 50 percent more callbacks than blacks. To address this concern, we conducted a separate experiment in which we only asked the main outcome questions in an obfuscated follow-up study one week after administering the treatment (Figure A.2 provides a summary of the structure).

Sample In Experiment 2, in collaboration with Research Now, one of the leading marketing research companies in the US, we recruited 2075 respondents for the first wave of the experiment in June 2017. The first wave was the second component of a follow-up study from another experiment that we also conducted with Research Now.¹⁵ Out of these 2075 respondents, 1720 also completed the second wave. As shown in Table A.2, the sample is broadly representative of the US population in terms of several important observable characteristics, such as race, household income, region of residence, and political affiliation. There is balance across treatment arms (Tables A.4 and A.5) and treatment status is not correlated with completing the follow-up (Table A.6).

Design The experiment had two waves. In the first wave, we elicited beliefs about racial discrimination in the same way as in Experiment 1. We then administered the same information treatment as in Experiment 1 to respondents in the treatment group. We intentionally did not ask any of the main outcome questions in the first wave to minimize

¹⁵ In the first wave, respondents also answered demographic questions, questions about their views on the role of the government, and questions about their views on immigration.

the risk that respondents would realize the connection between the two waves.

Approximately one week after the first wave, respondents were invited to participate in the second wave. We took several steps to obfuscate the purpose of the second wave. First, respondents received a generic invitation from the survey provider to participate in a five-minute survey which did not reveal that the two waves were connected (Figure A.12 provides a screenshot of the invitation from wave 1).¹⁶ Second, we used different Qualtrics accounts with different layouts for the two studies. Third, we asked respondents several obfuscation questions about their views on investment and religion before asking our main outcome questions.

Following the obfuscation questions, we asked the same questions on self-reported support for pro-black policies as in Experiment 1. We also asked some additional questions on potential mechanisms that were not included in Experiment 1, such as beliefs about the effectiveness of affirmative action policies. At the end of the survey, we elicited posterior beliefs about labor market discrimination using the same instructions and incentives as in the first wave. As we use the same belief elicitation across the two waves, it is natural to assume that respondents realized that the two waves are connected at this point.

Results: Posterior beliefs about labor market discrimination The treatment strongly weakens the correlation between prior and posterior beliefs about labor market discrimination (as shown in Figure A.6) and there is a convergence in beliefs between overestimators and underestimators (column 3 of Table 2). There is also a positive and significant interaction effect between the perception gap and the treatment: In response to learning that blacks had to send out one more resume than previously thought, treated respondents on average update their belief about the number of resumes that blacks needed to send out by 0.58 ($p < 0.01$; column 4 of Table 2). Consistent with the findings from the main experiment, there is no significant treatment heterogeneity based on political affiliation (columns 3 and 4 of Table A.8). Furthermore, treatment effects on posterior beliefs are stronger for respondents with less confidence in their pre-treatment beliefs (as shown in Table A.10),

¹⁶ The actual number of days between wave 1 and wave 2 varied between one and 19 days for all respondents, with an average of eight days.

consistent with genuine belief updating. Finally, using the results from Bertrand and Mullainathan (2004) as the benchmark, we find that treated respondents on average hold less biased beliefs about racial labor market discrimination compared to control group respondents (columns 4 to 6 of Table A.9).

Results: Self-reported support for pro-black policies As shown in columns 4 to 6 of Table A.14, there is an overall muted response to the information treatment on support for pro-black policies. If anything, there is a decrease in support for pro-black policies among underestimators. This “backfire” effect is almost entirely driven by Republican underestimators who reduce their support for pro-black policies by up to 30 percent of a standard deviation. While the effect is driven by a relatively small subsample and should therefore be interpreted with caution, a potential explanation for the backfire effect could be that the treatment simultaneously changes people’s beliefs about how effective affirmative action programs have been in helping blacks.¹⁷ As shown in column 1 of Table A.17, treated Republican underestimators are indeed more likely to think that affirmative action programs have hurt blacks ($p < 0.01$). Taken together, the evidence from Experiment 2 is consistent with the main finding from Experiment 1: while the information provision leads to a convergence in beliefs about the extent of racial discrimination, it does not lead to a convergence in support for pro-black policies.

B. Experiment 3: Robustness to framing and incentives

Motivation While our approach of eliciting quantitative and incentivized beliefs about the results of correspondence studies has many advantages compared to traditional survey questions, there are also some disadvantages. One potential concern is that respondents might believe that researchers are liberally biased. Incentives for accuracy then encourage

¹⁷ Another reason to interpret the backfire effect with caution is that we did not see any evidence of backfire effects in the main experiment. However, it could be the case that the backfire effect only arises in the obfuscated follow-up because Republican underestimators did not feel it was socially acceptable to express very low support for pro-black policies immediately after being informed by the experimenter that discrimination is more prevalent than their initial estimates.

respondents to report a combination of their true beliefs about discrimination and their beliefs about researcher bias. A second concern is that the results from the belief elicitation might depend on the question framing (Eriksson and Simpson, 2012). Finally, a concern with measuring belief updating about research results is that the treatment might differentially affect beliefs about whether researchers are liberally biased. To address these three concerns, we ran a third experiment with the following key design changes compared to the main experiment: (i) non-incentivized belief elicitations, (ii) three different ways of framing the pre-treatment belief elicitation, and (iii) a post-treatment belief elicitation that is not related to research results (and thus unrelated to beliefs about researcher bias).

Sample The experiment was conducted in June 2019 in collaboration with Lucid, a provider of representative online panels that is widely used in economic research (Burszтын et al., 2020). We recruited a sample of 2143 respondents broadly representative of the US population in terms of some important observable characteristics (age, income, region, and gender). Table A.2 provides summary statistics and Table A.7 shows that there is balance across treatment arms.

Design We randomized respondents into three different conditions with different ways of framing the pre-treatment belief elicitation about the results from the study by Bertrand and Mullainathan (2004). In contrast to the main experiment, we elicited these beliefs without any accuracy incentives. In the first condition, we used the same frame as in the main experiment; i.e., we elicit beliefs about the number of times a resume with a black-sounding name had to be sent out to get one callback on average after informing respondents about the true number for resumes with white-sounding names. In the second condition, we elicit beliefs about the number of times a resumes with *white-sounding* names had to be sent out to get one callback on average after informing respondents about the true number for resumes with black-sounding names. In the third condition, we first asked respondents whether resumes with white-sounding names or resumes with black-sounding names were more likely to receive callbacks for interviews. In a second step, we asked respondents to estimate how many percent more callbacks resumes with white-sounding or black-sounding names received (depending on their answer to the first

question).

After eliciting beliefs, we randomly assigned half of the respondents to an information treatment in which treated respondents were told that white-sounding names received 50 percent more callbacks for interviews than black-sounding names. We then measured self-reported support for pro-black policies using the same questions as in the main experiment. Finally, we measured non-incentivized post-treatment beliefs about the percent chance that there is racial discrimination in hiring without relating this belief elicitation to results from research studies.

Results: Robustness of prior beliefs to question framing We find broadly similar patterns based on the different ways of framing the belief elicitation. Across conditions, the largest fraction of our respondents report that there is racial discrimination against blacks, but the exact level differs somewhat between conditions (as shown in Figure A.4). Reassuringly, as shown in Panel A of Figure A.7, the correlates of beliefs about racial discrimination are very stable across the conditions. In particular, differences in beliefs between Republicans and non-Republicans are substantial and highly robust across the different conditions.

Results: Belief updating about racial discrimination in hiring We still see evidence of significant belief updating when beliefs are elicited without accuracy incentives in a general setting without reference to any academic studies. Specifically, treated underestimators increase their perception of the likelihood that there is hiring discrimination against blacks by 5.6 percentage points (column 5 of Table 2; $p < 0.01$). By contrast, we find no evidence that overestimators change their beliefs in response to the treatment.¹⁸ Examining political heterogeneity in treatment responses, we find that the effect on underestimators is mostly

¹⁸ While the belief elicitation in the main experiment measured quantitative beliefs about the extent of racial discrimination, the belief elicitation in this experiment only measures beliefs about whether there is discrimination against blacks. Given that all overestimators already held a qualitatively correct belief (namely that there is discrimination against blacks), it is theoretically ambiguous whether they should adjust their beliefs about whether there is discrimination against blacks in response to the information.

driven by Republican respondents (column 5 of Table A.8), which likely reflects that these respondents had more scope to change their beliefs about whether there was discrimination against blacks.

Results: Self-reported support for pro-black policies We find fairly comparable treatment effects across the different ways of framing the belief elicitation, but the effects are not very precisely estimated due to the large number of cells (Table A.16). Overall, as shown in columns 9-12 of Table A.14, the results are consistent with our main finding from Experiment 1 that self-reported policy views are largely unresponsive to information about racial discrimination. Furthermore, pooling results from Experiments 1, 2 and 3 (N=5,220), we observe precisely estimated null effects on self-reported policy views for both overestimators and underestimators (Table A.12). Taken together, the three information provision experiments clearly demonstrate that although providing information can substantially reduce disagreement about the extent of racial discrimination, it is not sufficient to reduce disagreement about pro-black policies

C. *Experiment 4: Beliefs about the sources of discrimination*

Motivation How beliefs about the extent of racial discrimination affect support for pro-black policies may critically depend on beliefs about the underlying sources of the discrimination. For instance, if people think that employers are primarily engaging in statistical discrimination, they might be less responsive to information about the extent of racial discrimination than if 1,720 We collect data to examine this potential mechanism in Experiment 4.¹⁹

Sample and design We conducted Experiment 4 in June 2019 and recruited, in collaboration with Lucid, a sample broadly representative of the US population in terms of gender, age, region, education, and income (N=1060). After explaining the design of the correspondence study by Bertrand and Mullainathan (2004) and the findings from the study, we asked our respondents what they think is the main reason for why employers are more likely to call back applicants with white-sounding names. We designed the possible

¹⁹ In Section E of the Online Appendix, we provide evidence that most people think that the callback differential hurts blacks in the labor market.

responses to match the most commonly cited theoretical reasons for differences in callback rates as closely as possible.²⁰

Results As shown in Figure A.10, the most common belief among our respondents is that the callback differential is due to implicit discrimination: 39 percent of our respondents think employers subconsciously rely on negative stereotypes about blacks. 17 percent think the main reason is inaccurate statistical discrimination (employers incorrectly believing that blacks are less productive on average), while 15 percent think the main reason is some form of accurate statistical discrimination (higher variance of unobserved skills or lower productivity on average). Only five percent of our respondents think the main reason is taste-based discrimination (i.e. that employers dislike interactions with blacks). The remaining 25 percent of our respondents do not think it is generally true that whites are more likely to get callbacks than blacks.

Conditional on believing that correspondence studies present evidence of discrimination, Republicans and Democrats alike tend to believe that discrimination is due to implicit discrimination or inaccurate statistical discrimination rather than taste-based or accurate statistical discrimination. However, 38 percent of Republicans say they do not think it is generally true that employers are more likely to call back white applicants compared to only 13 percent of Democrats. While we generally find that Republicans strongly update their beliefs in response to the research evidence, this result suggests that there is substantial heterogeneity among Republicans in how much they trust the research

²⁰ Specifically, we examine whether people think that lower callback rates are due to taste-based discrimination (Becker, 1957), accurate statistical discrimination due to lower average productivity of blacks (Arrow, 1973; Phelps, 1972), statistical discrimination due to blacks having a higher variance of unobserved skills (Aigner and Cain, 1977), inaccurate discrimination due to biases in beliefs (Bohren et al., 2019), or implicit discrimination due to subconscious negative stereotyping of blacks (Bertrand et al., 2005).

evidence.^{21,22}

D. Experiments 5 and 6: Are policy views elastic to other types of information?

Our main finding is that although providing information about results from correspondence studies can substantially reduce disagreement about the extent of racial discrimination, it is not sufficient to reduce disagreement about pro-black policies. This finding raises the question of whether support for pro-black policies is elastic to other types of information. In two additional experiments, we therefore explored whether other pieces of information affect support for pro-black policies.²³

²¹ We also asked whether respondents thought that sending out fictitious resumes to assess whether white-sounding names or black-sounding names receive more callbacks for interviews is a reliable method to detect racial discrimination in hiring. Although the majority of the respondents agree with this statement, there is substantial disagreement among Republicans in particular (Figure A.10). Consistent with fewer Republicans thinking correspondence studies are a good method to detect discrimination, incentivized evidence from Experiment 2 shows that Republicans display lower willingness to pay for the research evidence than non-Republicans (Table A.18). Furthermore, in Experiment 3, 51 percent of treated Republicans think it is more likely that blacks are discriminated against in the labor market, compared to 40 percent of Republicans in the control group. Even though a substantial fraction of treated Republicans thus remained unconvinced that blacks are discriminated against in the labor market, we still find evidence of strong updating among Republicans on average, consistent with substantial heterogeneity in the updating.

²² Differences in support for affirmative action across our different panels are unlikely to explain why we see evidence of belief updating in Experiments 1–3, including among Republicans, yet relatively high distrust in the research evidence among Republicans in Experiment 4. While we see some variation in support for affirmative action policies across our different data collections, as displayed in Figure A.11, we find that the Republicans from Experiment 4 display similar support for affirmative action as the Republicans in the main experiment.

²³ Appendix D describes the two experiments in more details.

Experiment 5 was motivated by strong correlational evidence suggesting an important role of beliefs about differences in work ethic between blacks and whites for explaining views on pro-black policies. In this experiment, conducted with 2999 participants recruited from Amazon Mechanical Turk, we provided our respondents with information challenging the stereotype that blacks have a worse work ethic than whites (Gilens, 2009). The experiment reveals that people who receive information about racial differences in work ethic do not adjust their views on pro-black policies.

Experiment 6 sheds light on a different prominently discussed causal determinant of policy views and political polarization, namely political identity (Bursztyn et al., 2019) and the importance of party cues (Brader and Tucker, 2012). In this experiment, conducted with 4000 participants recruited in collaboration with Research Now, we show that providing information about how the Republican and Democratic parties differ in their support for affirmative action does not affect Democrat–Republican differences in self-reported policy views.

Overall, these two additional experiments corroborate our finding that self-reported attitudes towards pro-black policies are generally hard to move with information, suggesting that these may have an important “cultural” component that is very stable over time (Luttmer and Singhal, 2011).

V. Concluding remarks

In this paper, we provide representative evidence on beliefs about racial discrimination and examine the scope for information to reduce disagreement in beliefs and support for pro-black policies. We document a large dispersion in beliefs about the extent of racial discrimination. Providing evidence from a correspondence study substantially reduces disagreement in beliefs but is largely ineffective in reducing disagreement over pro-black policies.

Our paper introduces a new approach for measuring beliefs about discrimination by eliciting priors about the results from correspondence studies. The main advantage of this approach is that it allows for the elicitation of quantitative beliefs that are easily comparable across respondents. Furthermore, this approach allows for the provision of

research evidence based on clean causal evidence. Our study demonstrates the feasibility of this approach by showing that correspondence studies can be easily explained to and understood by a general population sample. The approach could be useful for researchers who wish to study beliefs about discrimination against different groups, such as women.

Our findings have several implications for potential information campaigns about racial discrimination. If the goal of the information campaign is to reduce disagreement about policy issues, a broadly targeted information campaign about the extent of racial discrimination is unlikely to be very effective for two reasons. First, a large fraction of the population is unlikely to change their policy views even if their beliefs about racial discrimination change. Second, beliefs about racial discrimination are very dispersed, making it hard to predict the average impact of a campaign as people might update their beliefs in different directions. For a campaign to be successful, it is therefore important to target segments of the population that are both open to change their minds on policy and for which beliefs are not very dispersed. However, the overall likely muted effects of information campaigns on policy attitudes could be specific to the US context. Informational campaigns about racial discrimination might have more scope to change policy views in countries where people are more likely to underestimate the extent of racial discrimination relative to the information provided and where the political climate is less polarized.

References

- Aigner, Dennis J and Glen G Cain, “Statistical Theories of Discrimination in Labor Markets,” *ILR Review*, 1977, *30* (2), 175–187.
- Alesina, Alberto and Eliana La Ferrara, “A Test of Racial Bias in Capital Sentencing,” *American Economic Review*, 2014, *104* (11), 3397–3433.
- , Armando Miano, and Stefanie Stantcheva, “The polarization of reality,” in “AEA Papers and Proceedings,” Vol. 110 2020, pp. 324–28.
- , Michela Carlana, Eliana La Ferrara, and Paolo Pinotti, “Revealing Stereotypes: Evidence from Immigrants in Schools,” Technical Report, NBER Working Paper No. 25333 2018.
- , Stefanie Stantcheva, and Edoardo Teso, “Intergenerational Mobility and Preferences for Redistribution,” *American Economic Review*, 2018, *108* (2), 521–554.
- Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Y Yang, “Polarization and public health: Partisan differences in social distancing during the Coronavirus pandemic,” *NBER Working Paper*, 2020, (w26946).
- Arrow, Kenneth, “The Theory of Discrimination,” in Orley Ashenfelter and Albert Rees, eds., *Discrimination in the Labor Market*, Princeton University Press, 1973, pp. 3–33.
- Arrow, Kenneth J, “What Has Economics to Say about Racial Discrimination?,” *Journal of Economic Perspectives*, 1998, *12* (2), 91–100.
- Becker, Gary S, *The Economics of Discrimination*, University of Chicago Press, 1957.
- Bertrand, Marianne and Esther Duflo, “Field Experiments on Discrimination,” in Esther Duflo and Abhijit Banerjee, eds., *Handbook of Field Experiments*, Vol. 1, Elsevier, 2017, chapter 8.
- and Sendhil Mullainathan, “Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination,” *American Economic Review*, 2004, *94* (4), 991–1013.

- , Dolly Chugh, and Sendhil Mullainathan, “Implicit Discrimination,” *American Economic Review*, 2005, pp. 94–98.
- Bobo, Lawrence and James R Kluegel, “Opposition to Race-Targeting: Self-Interest, Stratification Ideology, or Racial Attitudes?,” *American Sociological Review*, 1993, pp. 443–464.
- Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G Pope, “Inaccurate Statistical Discrimination,” NBER Working Paper 25935 2019.
- Bond, Timothy N and Kevin Lang, “The Sad Truth About Happiness Scales,” *Journal of Political Economy*, 2019, 127 (4), 1629–1640.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, “Stereotypes,” *Quarterly Journal of Economics*, 2016, 131 (4), 1753–1794.
- Brader, Ted and Joshua A Tucker, “Following the Party’s Lead: Party Cues, Policy Opinion, and the Power of Partisanship in Three Multiparty Systems,” *Comparative Politics*, 2012, 44 (4), 403–403.
- Bullock, John G, Alan S Gerber, Seth J Hill, and Gregory A Huber, “Partisan bias in factual beliefs about politics,” *Quarterly Journal of Political Science*, 2015, 10, 519–578.
- Bursztyjn, Leonardo, Ingar Haaland, Aakaash Rao, and Christopher Roth, “Disguising Prejudice: Popular Rationales as Excuses for Intolerant Expression,” NBER Working Paper 27288 May 2020.
- , Michael Callen, Bruno Ferman, Saad Gulzar, Ali Hasanain, and Noam Yuchtman, “Political Identity: Experimental Evidence on Anti-Americanism in Pakistan,” *Journal of the European Economic Association*, 09 2019. jvz053.
- Cappelen, Alexander, Ingar Haaland, and Bertil Tungodden, “Beliefs about Behavioral Responses to Taxation,” *Working Paper*, 2017.

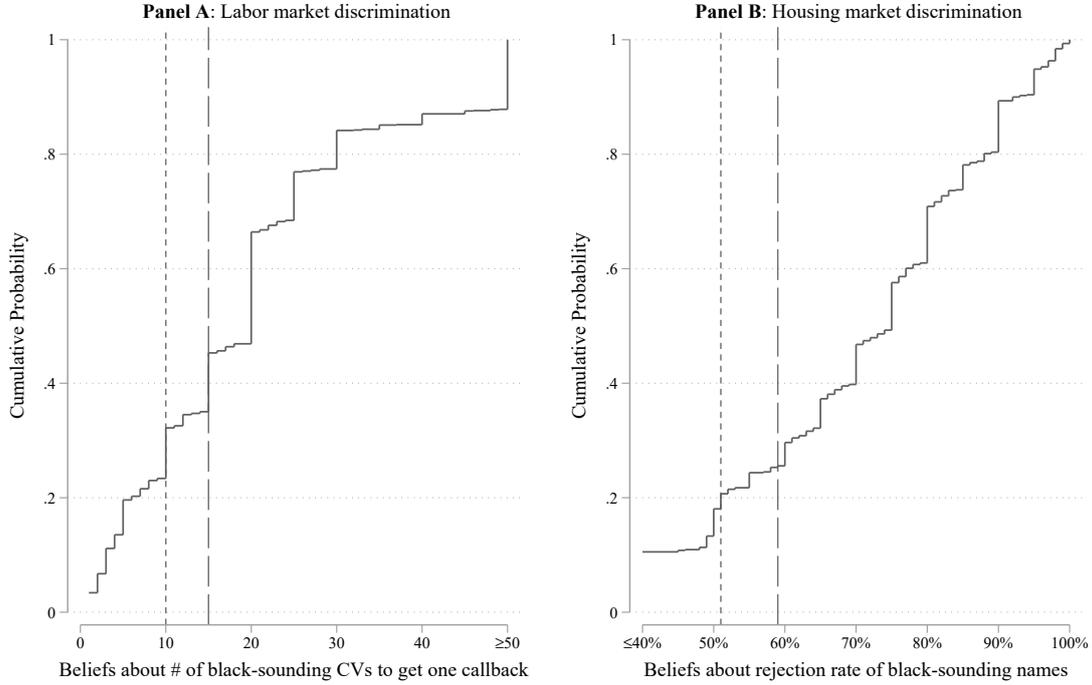
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia, “Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments,” *American Economic Journal: Macroeconomics*, 2016, 9 (3), 1–35.
- Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz, “Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment,” *Journal of Public Economics*, 2013, 98, 100–112.
- de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth, “Measuring and Bounding Experimenter Demand,” *American Economic Review*, 2018, 108 (11), 3266–3302.
- DellaVigna, Stefano and Devin Pope, “Predicting Experimental Results: Who Knows What?,” *Journal of Political Economy*, 2018, 126 (6), 2410–2456.
- and —, “What Motivates Effort? Evidence and Expert Forecasts,” *Review of Economic Studies*, 2018, 85 (2), 1029–1069.
- , —, and Eva Vivalt, “Predict science to improve science,” *Science*, 2019, 366 (6464), 428–429.
- Edelman, Benjamin G, Michael Luca, and Dan Svirsky, “Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment,” *American Economic Journal: Applied Economics*, 2017, 2, 1–22.
- Eriksson, Kimmo and Brent Simpson, “What do Americans know about inequality? It depends on how you ask them,” *Judgment and Decision Making*, 2012, 7, 741–745.
- Gilens, Martin, *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy*, University of Chicago Press, 2009.
- Grigorieff, Alexis, Christopher Roth, and Diego Ubfal, “Does Information Change Attitudes Toward Immigrants?,” *Demography*, 2020, 57 (3), 1–27.
- Haaland, Ingar and Christopher Roth, “Labor market concerns and support for immigration,” *Journal of Public Economics*, 2020, 191, 104256.

- , —, and Johannes Wohlfart, “Designing Information Provision Experiments,” Working paper 8406, CESifo June 2020.
- Harrison, David A, David A Kravitz, David M Mayer, Lisa M Leslie, and Dalit Lev-Arey, “Understanding attitudes toward affirmative action programs in employment: Summary and meta-analysis of 35 years of research.,” *Journal of Applied Psychology*, 2006, *91* (5), 1013–36.
- Heckman, James J, “Detecting Discrimination,” *Journal of Economic Perspectives*, 1998, *12* (2), 101–116.
- Hjort, Jonas, Diana Moreira, Gautam Rao, and Juan Francisco Santini, “How Research Affects Policy: Experimental Evidence from 2,150 Brazilian Municipalities,” *American Economic Review*, *forthcoming*, 2020.
- Horton, John J, David G Rand, and Richard J Zeckhauser, “The Online Laboratory: Conducting Experiments in a Real Labor Market,” *Experimental Economics*, 2011, *14* (3), 399–425.
- Karadja, Mounir, Johanna Mollerstrom, and David Seim, “Richer (and Holier) Than Thou? The Effect of Relative Income Improvements on Demand for Redistribution,” *Review of Economics and Statistics*, 2017, *99* (2), 201–212.
- Kluegel, James R and Eliot R Smith, “Affirmative action attitudes: Effects of self-interest, racial affect, and stratification beliefs on Whites’ views,” *Social Forces*, 1983, *61* (3), 797–824.
- Kuklinski, James H, Paul M Sniderman, Kathleen Knight, Thomas Piazza, Philip E Tetlock, Gordon R Lawrence, and Barbara Mellers, “Racial Prejudice and Attitudes Toward Affirmative Action,” *American Journal of Political Science*, 1997, *4* (2), 402–419.
- Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva, “How Elastic are Preferences for Redistribution? Evidence from Randomized Survey Experiments,” *American Economic Review*, 2015, *105* (4), 1478–1508.

- List, John A, “The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field,” *Quarterly Journal of Economics*, 2004, 119 (1), 49–89.
- Luttmer, Erzo FP and Monica Singhal, “Culture, Context, and the Taste for Redistribution,” *American Economic Journal: Economic Policy*, 2011, 3 (1), 157–79.
- Newkirk, Vann R., “*The Myth of Reverse Racism*,” 2017. <https://www.theatlantic.com/education/archive/2017/08/myth-of-reverse-racism/535689/> [Accessed: 16/11/2018].
- Phelps, Edmund S, “The Statistical Theory of Racism and Sexism,” *American Economic Review*, 1972, 62 (4), 659–661.
- Pope, Devin G, Joseph Price, and Justin Wolfers, “Awareness Reduces Racial Bias,” *Management Science*, 2018, 64 (11), 4967–5460.
- Prior, Markus, Gaurav Sood, and Kabir Khanna, “You Cannot be Serious: The Impact of Accuracy Incentives on Partisan Bias in Reports of Economic Perceptions,” *Quarterly Journal of Political Science*, 2015, 10, 489–518.
- Quillian, Lincoln, Devah Pager, Ole Hexel, and Arnfinn H Midtbøen, “Meta-analysis of field experiments shows no change in racial discrimination in hiring over time,” *Proceedings of the National Academy of Sciences*, 2017, p. 201706255.
- Smith, Tom W., Michael Davern, Jeremy Freese, and Michael Hout, “GSS General Social Survey,” Data accessed November 16, 2018 from the GSS Data Explorer website at gssdataexplorer.norc.org 2014.
- Taber, Charles S and Milton Lodge, “Motivated skepticism in the evaluation of political beliefs,” *American Journal of Political Science*, 2006, 50 (3), 755–769.
- Vivalt, Eva and Aidan Coville, “How Do Policymakers Update Their Beliefs?,” 2020.

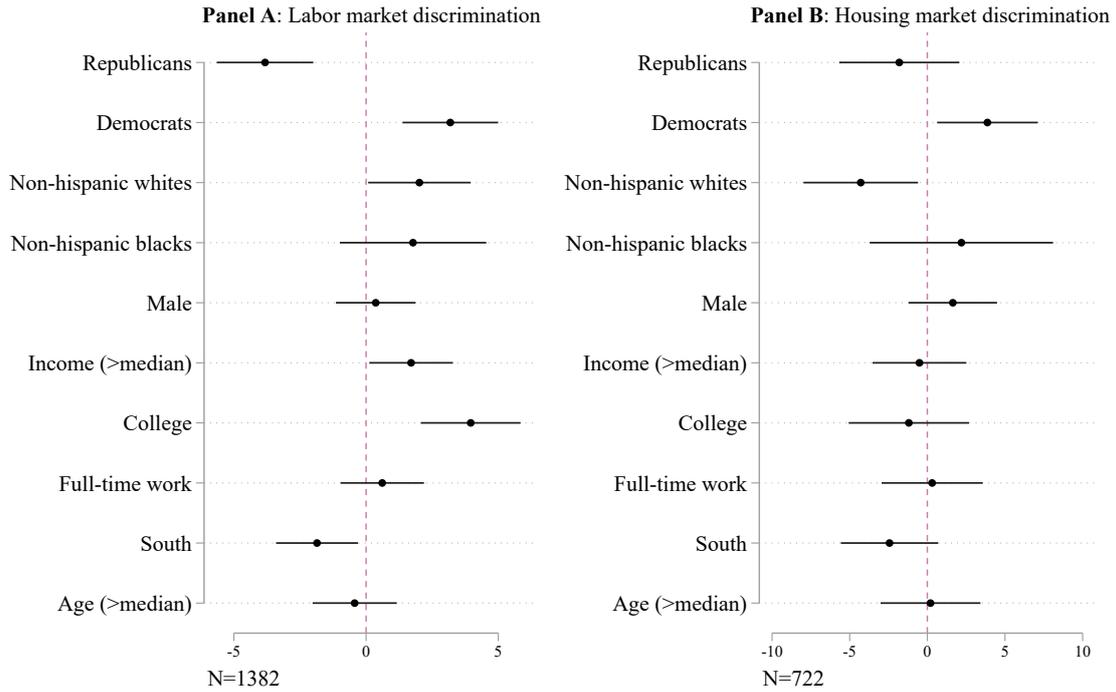
Figures and tables

Figure 1: Beliefs about racial discrimination in the labor and housing market



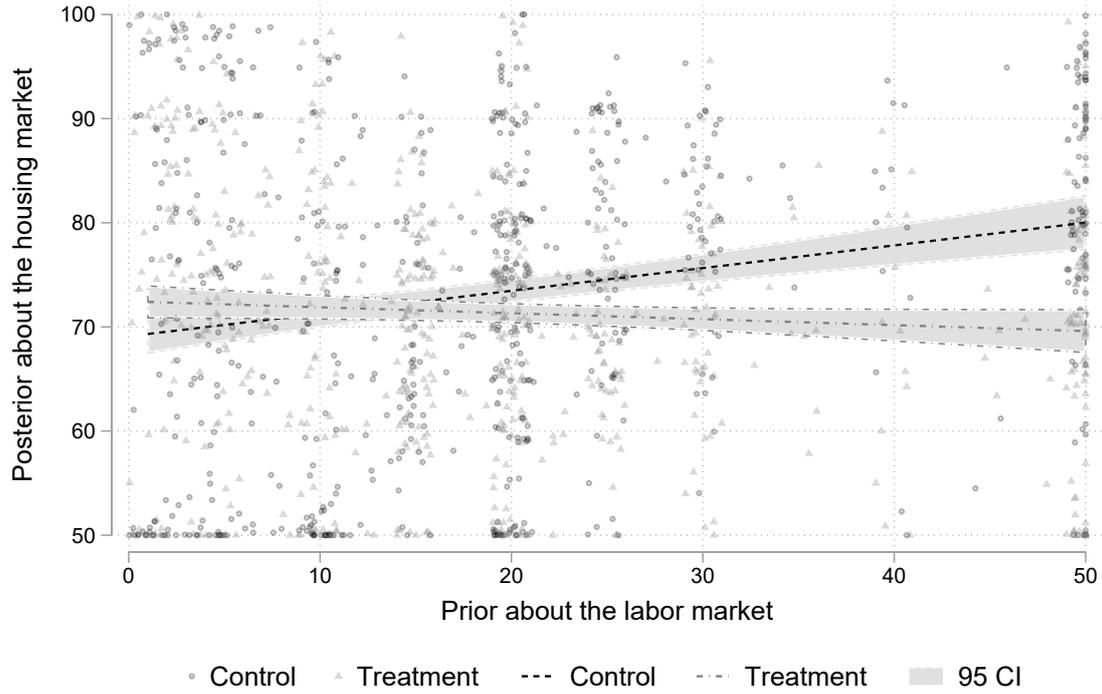
Notes: This figure uses data from Experiment 1 (the NORC sample). **Panel A** shows data on beliefs about how many times resumes with black-sounding names on average had to be sent out to get one callback for an interview. Respondents were informed that the corresponding number for resumes with white-sounding names was ten (as found in the study by Bertrand and Mullainathan, 2004). **Panel B**, using only control group respondents, shows data on beliefs about the rejection rate on reservation requests sent from accounts with black-sounding names. Respondents were initially asked about the percent rate of acceptances of reservation requests for black-sounding names on Airbnb (true rate is 41 percent, as found in the study by Edelman et al., 2017). They were told that the corresponding number for white-sounding names was 49. We have recoded the values to implied rejection rates by subtracting each estimate from 100. In both panels, the **short-dashed** lines indicate the true level for whites and the **long-dashed** lines indicate the true level for blacks.

Figure 2: Correlates of beliefs about racial discrimination



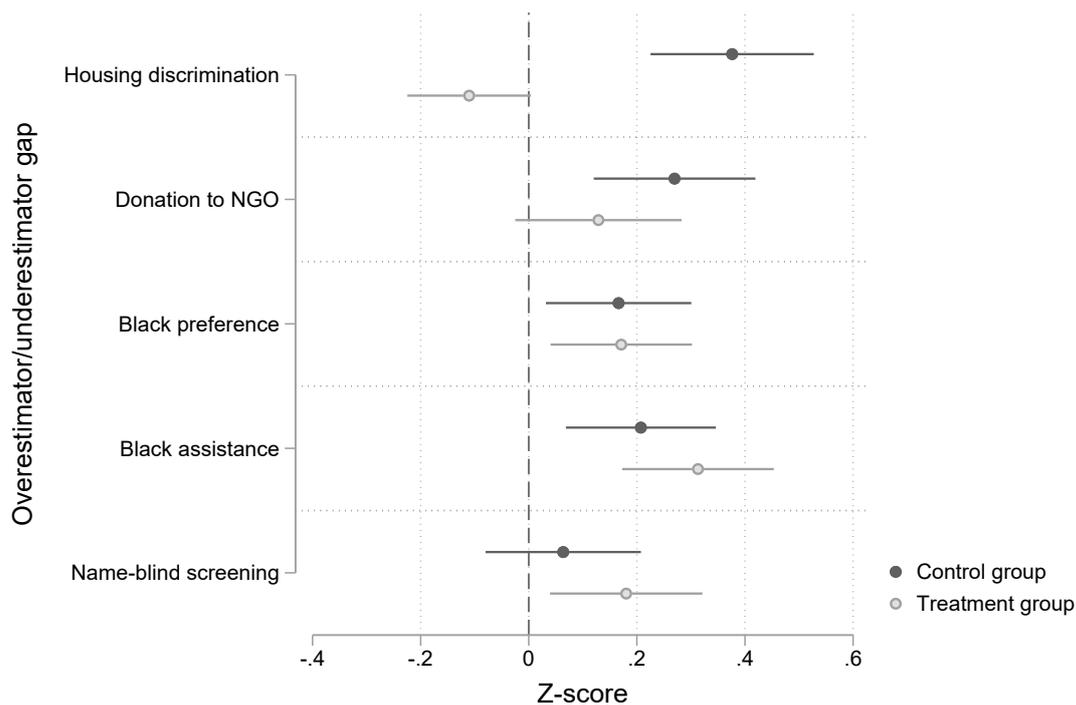
Notes: This figure uses data from Experiment 1 (the NORC sample). The dots indicate the mean values of the estimated multiple regression coefficients. The dependent variable in **Panel A** is people's beliefs about the number of times resumes with black-sounding names had to be sent out to get one callback for an interview on average. The dependent variable in **Panel B** is people's beliefs about the percent of time reservation requests from black-sounding names on Airbnb were rejected. Lines indicate 95 percent confidence intervals.

Figure 3: Belief updating in response to the research evidence



Notes: This figure, which uses data from Experiment 1 (NORC), shows a scatterplot with prior beliefs about racial discrimination in hiring on the x-axis (beliefs about the number of resumes with black-sounding names that had to be sent out to get one callback on average) and post-treatment beliefs about racial discrimination in the housing market on the y-axis (beliefs about the rejection rate of black-sounding names in the housing market) separately for the treatment and the control group. Prior (post-treatment) beliefs have been top (bottom) coded at 50. Lines indicate fitted values with 95 percent confidence intervals.

Figure 4: Overestimator–underestimator differences in beliefs and preferences: Treatment vs. control group



Notes: This figure uses data from Experiment 1 (the NORC sample). The dots indicate the z-scored difference in outcome variables (beliefs, donations, self-reported policy views) between those who initially overestimate the extent of racial discrimination in the labor market and those who initially underestimate the extent of racial discrimination. The estimates are obtained from multiple regressions that include pre-specified control variables run separately for treatment and control group respondents. Lines indicate 95 percent confidence intervals.

Table 1: The association between beliefs and preferences

	(1) Donations to NGO	(2) Black preference	(3) Black assistance	(4) Disc. housing	(5) Disc. ser. problem
Panel A: Without controls					
Beliefs about discrimination	0.219*** (0.040)	0.241*** (0.036)	0.246*** (0.035)	0.217*** (0.039)	0.294*** (0.035)
Panel B: With controls					
Beliefs about discrimination	0.171*** (0.041)	0.167*** (0.034)	0.169*** (0.035)	0.213*** (0.040)	0.231*** (0.031)
N	653	676	677	673	679

Note: This table show OLS regressions from control group respondents in Experiment 1 (NORC). In **Panel A**, we regress the outcome indicated in each column on standardized beliefs about racial discrimination in the labor market (i.e., beliefs about the number of times resumes with black-sounding names had to be sent out to receive one callback on average). In **Panel B**, we also include pre-specified controls in the regressions (gender, age, race, region, income, education, employment, and political views). *Donations to the NGO* refers to the number of times the respondents preferred to donate \$5 to the pro-black civil rights organization over money for themselves (responses range from 0 to 6). For the outcomes *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job) and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). *Disc. housing* refers to beliefs about the rejection rate of black-sounding names in the housing market (elicited on a scale from 0 to 100). *Disc. ser. problem* refers to the question of whether “racial discrimination against blacks in the labor market is a serious problem” which was elicited on a scale from 1 (Not a problem at all) to 5 (A very serious problem). All outcomes are z-scored.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 2: Belief updating

	Housing market (NORC)		Labor market (RN)		Labor market (Lucid)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (a)	4.16*** (1.54)	0.53 (1.05)	2.08** (1.02)	-2.51*** (0.75)	5.75*** (1.73)	4.17*** (1.22)
Overestimate \times Treatment (b)	-9.91*** (1.90)		-13.06*** (1.63)		-4.86** (2.34)	
Overestimate	7.61*** (1.53)		14.00*** (1.34)		14.07*** (1.66)	
Perception gap \times Treatment		0.37*** (0.07)		0.58*** (0.07)		0.14* (0.08)
Perception gap		-0.30*** (0.05)		-0.62*** (0.06)		-0.52*** (0.06)
Linear combination: a + b	-5.75*** (1.12)		-10.99*** (1.26)		0.88 (1.59)	
N	1366	1366	1701	1701	2098	2098
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean: Dependent variable	71.1	71.1	19.3	19.3	52.1	52.1

Note: The table shows OLS regression results where the dependent variables are post-treatment beliefs about what percent of the time reservation requests from black-sounding names were rejected on Airbnb (columns 1–2; Experiment 1 with NORC), post-treatment beliefs about the number of resumes with black-sounding names that had to be sent out to get one callback on average (columns 3–4; wave 2 of Experiment 2 with Research Now), and post-treatment beliefs about the percent chance that there is racial discrimination in hiring decisions against blacks (columns 5–6; Experiment 3 with Lucid). We include pre-specified controls (including gender, age, race, region, income, education, employment, and political views) in all specifications. For post-treatment beliefs about the labor market (columns 3 and 4), we also include confidence in prior beliefs as a control. “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market (i.e., who thought pre-treatment that resumes with black-sounding names had to send out more than 15 resumes to get one callback on average). The variable “Perception gap” is defined as Info–Prior, i.e. 15 minus the pre-treatment estimate of the number of resumes with black-sounding names that had to be sent out to get one callback on average.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table 3: Treatment effects: Main experiment

	(1) Housing discrimination	(2) Donation to NGO	(3) Name-blind screening	(4) Black preference	(5) Black assistance
Treatment (a)	0.211*** (0.078)	0.159** (0.075)	0.011 (0.076)	-0.028 (0.070)	-0.015 (0.077)
Treatment \times Overestimate (b)	-0.502*** (0.096)	-0.139 (0.107)	0.079 (0.101)	-0.037 (0.094)	0.059 (0.099)
Overestimate	0.386*** (0.078)	0.269*** (0.075)	0.086 (0.073)	0.194*** (0.068)	0.234*** (0.071)
Linear combination: a + b	-0.29*** (0.06)	0.02 (0.08)	0.09 (0.07)	-0.07 (0.06)	0.04 (0.06)
N	1366	1327	1378	1377	1374
Controls	Yes	Yes	Yes	Yes	Yes

Note: The table shows OLS regression results using respondents from Experiment 1 (NORC). *Housing discrimination* refers to beliefs about racial discrimination in the housing market (higher values indicate more discrimination). *Donations to the NGO* refers to the number of times the respondents preferred to donate \$5 to the pro-black civil rights organization over money for themselves (responses range from 0 to 6). For the outcomes *Name-blind screening* (support for mandatory name-blind recruitment), *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job), and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). All outcomes are z-scored using the mean and standard deviation in the control group. “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market. We include the following controls in all specifications: gender, age, race (indicators for blacks and whites), region (three indicators), income, education (indicator for having at least a two-year college degree), employment (indicator for having full-time work), and self-reported political affiliation (indicators for Republicans and Democrats).

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Online Appendix:

Beliefs About Racial Discrimination and Support for Pro-Black Policies

Ingar Haaland and Christopher Roth

Summary of the Online Appendix

Section A provides all of the appendix tables. Section A.1 provides an overview of all experiments (Table A.1), summary statistics for Experiments 1–3 (Table A.2), evidence of covariate balance for Experiments 1 and 2 (Table A.3, Table A.4, and Table A.5), results on attrition for the obfuscated follow-up study (Table A.6), and evidence of covariate balance for Experiment 3 (Table A.7). Section A.2 provides additional results on belief updating: Table A.8 shows political heterogeneity; Table A.9 shows treatment effects on accuracy (relative to our benchmark studies); and Table A.10 shows heterogeneity by confidence in prior beliefs. Section A.3 provides additional results on support for pro-black policies: Table A.11 shows results on donations, including political heterogeneity; Table A.12 shows results on self-reported support for-pro black policies when pooling results from the three information provision experiments; Table A.13 shows political heterogeneity using the pooled data; and Table A.14 shows results separately for each experiment, including political heterogeneity. Section A.4 provides additional results on robustness and mechanisms: Table A.15 and Table A.16 provide robustness results on elicitation techniques; Table A.17 provides treatment effects on some mechanism questions; and Table A.18 provides results on willingness to pay for the research evidence.

Section B provides all the appendix figures. Figure A.1 illustrates political differences in beliefs and preferences. Figure A.2 provides an overview of Experiments 1–3. Figure A.3 shows differences in beliefs between Republicans and Democrats. Figure A.4 shows descriptives across elicitation techniques. Figure A.5 shows political differences in donations. Figure A.6 shows belief updating in the labor market. Figure A.7 shows

correlates of prior beliefs across elicitation techniques. Figure A.8 shows correlates of beliefs about correspondence studies. Figure A.9 shows correlates of attitudes towards pro-black policies. Figure A.10 shows results on mechanisms, including beliefs about the sources of discrimination. Figure A.11 shows support for affirmative action across all experiments.

Section C provides screenshots of the consent forms for Experiment 2 and the recruitment email from Research Now. Section D describes Experiments 5 and 6 in more detail. Section E provides evidence on beliefs about the consequences of lower callback rates. Finally, Section F provides experimental instructions for all the experiments.

A. Appendix tables

A.1. Overview, summary statistics, balance and attrition

Table A.1: Overview of experiments

Experiment	Sample	Treatments Arms	Main outcomes
Pre-analysis plans			
Experiment 1 (June and July 2017)	NORC (N=1,382)	Treatment: Information about results from the correspondence study Control: No information	Donations to NGO; incentivized post-treatment beliefs; self-reported policy views
Pre-analysis Plan II			
Experiment 2: Wave 1 (June 2017)	Research Now (N=2,073)	Treatment: Information about results from the correspondence study Control: No information	None (elicited in wave 2)
Pre-analysis Plan I			
Experiment 2: Wave 2 (June and July 2017)	Research Now (N=1,720)	No treatments (administered in wave 1)	Incentivized post-treatment beliefs; self-reported policy views
Pre-analysis Plan I			
Experiment 3 (June 2019)	Lucid (N=2,143)	Treatment: Information about results from the correspondence study Control: No information	Non-incentivized beliefs and self-reported policy views
Pre-analysis Plan V			
Experiment 4 (June 2019)	Lucid (N=1,060)	No treatments	Beliefs about the reasons for differences in callback rates; self-reported policy views
Pre-analysis Plan VI			
Experiment 5 (October 2018)	MTurk (N=2,999)	Treatment: Information about inaccurate racial stereotypes Control: No information	Self-reported policy views
Pre-analysis Plan IV			
Experiment 6 (July 2018)	Research Now (N=4,000)	Treatment: Information about party views on affirmative action Control: No information	Self-reported policy views
Pre-analysis Plan III			

Notes: This table provides an overview of the different experiments conducted.

Table A.2: Summary statistics

	(1)	(2)	(3)
	Experiment 1	Experiment 2	Experiment 3
Male	0.465	0.506	0.493
Age	48.2	47.9	45.1
Household income	68157	73461	76989
African American/Black	0.112	0.062	0.089
Non-Hispanic white	0.672	0.485	0.779
College (at least 2-year degree)	0.807	0.819	0.744
Full-time employee	0.627	0.599	0.414
Republican	0.230	0.260	0.345
Democrat	0.358	0.383	0.344
Northeast	0.158	0.233	0.186
West	0.221	0.237	0.204
Midwest	0.295	0.187	0.199
South	0.326	0.344	0.411
Overestimate discrimination	0.547	0.460	0.577
Observations	1382	1720	2143

Note: This table displays summary statistics for the three main experiments: Experiment 1 with NORC, Experiment 2 with Research Now (the obfuscated follow-up survey), and Experiment 3 with Lucid.

Table A.3: Balance: Experiment 1 (NORC)

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Respondent age	49.06	47.27	0.047	1382
Male	0.45	0.48	0.156	1382
Non-Hispanic black	0.11	0.11	0.916	1382
Non-Hispanic white	0.68	0.66	0.508	1382
Northeast	0.16	0.15	0.721	1382
Midwest	0.27	0.32	0.034	1382
South	0.33	0.32	0.668	1382
West	0.24	0.20	0.127	1382
Household size	2.65	2.73	0.297	1382
Log household income	10.86	10.82	0.298	1382
At least some college	0.83	0.78	0.025	1382
Paid employee	0.53	0.52	0.851	1382
Self-employed	0.10	0.11	0.655	1382
Overestimate	0.54	0.55	0.708	1382
Prior (continuous, top-coded at 50)	20.53	19.72	0.300	1382
Republican	0.23	0.23	0.969	1382
Democrat	0.36	0.35	0.744	1382

Notes: This table displays covariate means for the treatment and control group for Experiment 1 (NORC). “Prior (dummy)” takes the value one for respondents who overestimate racial discrimination in the labor market. The p-value of a joint F-test of a regression of the treatment indicator on all of the covariates is $p=0.164$.

Table A.4: Balance: Experiment 2 (Research Now; baseline survey)

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Respondent age	47.19	47.66	0.493	2073
Male	0.50	0.49	0.844	2073
Non-Hispanic black	0.06	0.05	0.335	2073
Non-Hispanic white	0.49	0.48	0.812	2073
Household size	2.42	2.50	0.228	2073
Log household income	10.92	10.94	0.691	2073
At least 2-year college degree	0.83	0.82	0.609	2073
Overestimate	0.47	0.45	0.350	2073
Confidence in prior	3.31	3.36	0.295	2073
Republican	0.25	0.26	0.643	2073
Democrat	0.38	0.37	0.799	2073
West	0.22	0.24	0.225	2073
South	0.35	0.35	0.947	2073
Northeast	0.24	0.22	0.281	2073
Midwest	0.19	0.19	0.940	2073

Notes: This table displays covariate means for the treatment and control group (wave 1 of Experiment 2 with Research Now). “Prior (dummy)” takes the value one for respondents who overestimate racial discrimination in the labor market. “Confidence in prior” (i.e., confidence in the answer to the question of how many times resumes with black-sounding names had to be sent out to get one callback on average) was elicited on a scale from 1 (Very unsure) to 5 (Very Sure). The p-value of a joint F-test of a regression of the treatment indicator on all of the covariates is p=0.918.

Table A.5: Balance: Experiment 2 (Research Now; obfuscated follow-up)

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Respondent age	47.53	48.33	0.284	1720
Male	0.51	0.50	0.759	1720
Non-Hispanic black	0.07	0.06	0.502	1720
Non-Hispanic white	0.49	0.48	0.695	1720
Household size	2.43	2.45	0.751	1720
Log household income	10.92	10.94	0.703	1720
At least 2-year college degree	0.82	0.82	0.944	1720
Overestimate	0.47	0.45	0.422	1720
Prior (continuous, top-coded at 50)	17.06	17.13	0.916	1720
Confidence in prior	3.31	3.37	0.221	1720
Republican	0.25	0.27	0.569	1720
Democrat	0.39	0.38	0.730	1720
West	0.23	0.25	0.286	1720
South	0.34	0.35	0.786	1720
Northeast	0.24	0.22	0.306	1720
Midwest	0.19	0.18	0.701	1720

Notes: This table displays covariate means for the treatment and control group (wave 2 of Experiment 2 with Research Now). “Prior (dummy)” takes the value one for respondents who overestimate racial discrimination in the labor market. “Confidence in prior” (i.e., confidence in the answer to the question of how many times resumes with black-sounding names had to be sent out to get one callback on average) was elicited on a scale from 1 (Very unsure) to 5 (Very Sure). The p-value of a joint F-test of a regression of the treatment indicator on all of the covariates is p=0.961.

Table A.6: Experiment 2: Correlates of attrition

	Completed Follow-up	
	(1)	(2)
Treatment	-0.025 (0.017)	-0.027 (0.017)
Republican		0.049** (0.023)
Independent		0.041** (0.021)
Log(Income)		-0.001 (0.012)
College		-0.051** (0.024)
Black		0.036 (0.036)
White		-0.007 (0.019)
Prior (dummy)		0.016 (0.018)
Confidence in Prior		0.005 (0.009)
Male		0.042** (0.018)
Age		0.001 (0.001)
Response rate	0.806	0.806
Observations	2073	2073

Notes: The outcome variables takes value one if our respondent completed the follow-up study (wave 2 of Experiment 2 with Research Now). “Treatment” takes value one if the respondent received information about the results from the correspondence study. “Prior (dummy)” takes the value one for respondents who overestimate racial discrimination in the labor market. “Confidence in prior” (i.e., confidence in the answer to the question of how many times resumes with black-sounding names had to be sent out to get one callback on average) was elicited on a scale from 1 (Very unsure) to 5 (Very sure). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.7: Balance: Experiment 3 (Lucid)

	Treatment (T)	Control (C)	P-value(T - C)	Observations
Respondent age (cont.)	44.71	45.52	0.235	2143
Male	0.48	0.51	0.235	2143
African American/Black	0.09	0.09	0.956	2143
Non-Hispanic White	0.78	0.78	0.781	2143
Log household income	10.94	10.99	0.160	2143
At least some college	0.74	0.75	0.613	2143
Overestimate discrimination	0.57	0.59	0.342	2143
Republican	0.34	0.35	0.715	2143
Democrat	0.34	0.35	0.784	2143
West	0.20	0.21	0.319	2143
South	0.41	0.41	0.973	2143
Northeast	0.19	0.18	0.721	2143
Midwest	0.20	0.19	0.537	2143

Notes: This table displays covariate means for the treatment and control group (Lucid). “Prior (dummy)” takes the value one for respondents who overestimate racial discrimination in the labor market. The p-value of a joint F-test of a regression of the treatment indicator on all of the covariates is $p=0.461$.

A.2. Additional results on belief updating

Table A.8: Belief updating: Heterogeneity by political views

	Housing market (NORC)		Labor market (RN)		Labor market (Lucid)	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Non-Republicans						
Treatment (a)	2.84 (1.84)	-0.65 (1.23)	1.60 (1.21)	-2.64*** (0.87)	1.62 (2.25)	1.99 (1.60)
Overestimate × Treatment (b)	-9.34*** (2.20)		-12.96*** (1.90)		-1.24 (2.95)	
Overestimate	6.73*** (1.76)		14.20*** (1.58)		15.12*** (2.12)	
Perception gap × Treatment		0.33*** (0.08)		0.62*** (0.08)		0.10 (0.10)
Perception gap		-0.27*** (0.06)		-0.66*** (0.07)		-0.52*** (0.07)
Linear combination: a + b	-6.50*** (1.21)		-11.36*** (1.45)		0.39 (1.91)	
N	1052	1052	1255	1255	1375	1375
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean: Dependent variable	72.3	72.3	20.1	20.1	59.3	59.3
Panel B: Republicans						
Treatment (a)	7.55*** (2.88)	3.73* (2.06)	3.28* (1.91)	-1.62 (1.48)	11.15*** (2.62)	7.45*** (1.89)
Overestimate × Treatment (b)	-10.60** (4.22)		-12.65*** (3.27)		-8.97** (3.86)	
Overestimate	9.84*** (3.34)		12.75*** (2.61)		10.87*** (2.64)	
Perception gap × Treatment		0.48** (0.19)		0.41*** (0.13)		0.09 (0.15)
Perception gap		-0.40*** (0.12)		-0.43*** (0.11)		-0.46*** (0.10)
Linear combination: a + b	-3.05 (2.91)		-9.37*** (2.55)		2.18 (2.81)	
N	314	314	446	446	723	723
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean: Dependent variable	67.2	67.2	17.3	17.3	38.7	38.7

Note: The table shows OLS regression results where the dependent variables are post-treatment beliefs about what percent of the time reservation requests from black-sounding names were rejected on Airbnb (columns 1–2; Experiment 1 with NORC) and post-treatment beliefs about the number of resumes with black-sounding names on average had to be sent out to get one callback on average (columns 3–4; wave 2 of Experiment 2 with Research Now). In even-numbered columns, we include pre-specified controls (including gender, age, race, region, income, education, employment, and political views). “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market (i.e., who thought pre-treatment that resumes with black-sounding names had to send out more than 15 resumes to get one callback on average). The variable “Perception gap” is defined as Info–Prior, i.e. 15 minus the pre-treatment estimate of the number of resumes with black-sounding names that had to be sent out to get one callback on average.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.9: Treatment effects on biases in beliefs about racial discrimination

	Labor market (RN)			Housing market (NORC)		
	(1) Full sample	(2) Non-Republicans	(3) Republicans	(4) Full sample	(5) Non-Republicans	(6) Republicans
Panel A: Average effect						
Treatment	-6.66*** (0.67)	-7.45*** (0.80)	-4.44*** (1.27)	-3.93*** (0.63)	-4.47*** (0.69)	-2.12 (1.41)
N	1701	1255	446	1366	1052	314
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean: Dependent variable	13.3	13.8	11.9	19.4	20.1	17.4
Panel B: Heterogeneity by prior						
Treatment	-2.22*** (0.80)	-2.76*** (0.95)	-0.93 (1.49)	-2.31** (1.03)	-3.21*** (1.22)	-0.19 (1.91)
Overestimate \times Treatment	-9.64*** (1.37)	-9.70*** (1.59)	-8.87*** (2.71)	-2.96** (1.28)	-2.18 (1.46)	-4.55 (2.82)
N	1701	1255	446	1366	1052	314
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Note: The table shows OLS regression results where the dependent variables are the absolute value of post-treatment biases in beliefs about the about the number of resumes with black-sounding names that had to be sent out to get one callback on average (columns 1–3; wave 2 of Experiment 2 with Research Now) and the absolute value of post-treatment biases in beliefs about what percent of the time reservation requests from black-sounding names were rejected on Airbnb (columns 4–6; Experiment 1 with NORC). To calculate the bias, we subtract the correct answer (as identified in the studies by Bertrand and Mullainathan (2004) and Edelman et al. (2017) for the labor and housing market, respectively) from people’s estimates. We include pre-specified controls (including gender, age, race, region, income, education, employment, and political views) in all specifications. “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.10: Belief updating: Heterogeneity by confidence in prior beliefs

	Posterior: Labor market	
	(1)	(2)
Treatment	11.88*** (4.11)	3.95 (2.95)
Confidence \times Treatment	-2.88** (1.20)	-1.91** (0.88)
Overestimate \times Treatment	-23.29*** (6.12)	
Perception gap \times Treatment		0.85*** (0.26)
Overestimate \times Confidence \times Treatment	3.00* (1.82)	
Perception gap \times Confidence \times Treatment		-0.08 (0.08)
Overestimate	19.00*** (5.02)	
Perception gap		-0.75*** (0.22)
Confidence	1.36 (0.94)	0.97 (0.70)
N	1701	1701
Controls	Yes	Yes

Note: The table shows OLS regression results where the dependent variable is post-treatment beliefs about the number of resumes with black-sounding names had to be sent out to get one callback on average (wave 2 of Experiment 2 with Research Now). In column 2, we include pre-specified controls (including gender, age, race, region, income, education, employment, and political views). “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market (i.e., who thought pre-treatment that resumes with black-sounding names had to be sent out more than 15 times to get one callback on average). The variable “Perception gap” is defined as Info–Prior, i.e. 15 minus the pre-treatment estimate of the number of resumes with black-sounding names that had to be sent out to get one callback on average. “Confidence” refers to confidence in pre-treatment beliefs (measured instantly after the belief elicitation), which was elicited on a scale from 1 (Very unsure) to 5 (Very sure).

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

A.3. Additional results on support for pro-black policies

Table A.11: Treatment effects on donations

	Full sample		Non-Republicans		Republicans	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (a)	0.159** (0.075)	0.118** (0.055)	0.206** (0.093)	0.135** (0.066)	0.070 (0.123)	0.083 (0.100)
Treatment \times Overestimate (b)	-0.139 (0.107)		-0.215* (0.125)		0.015 (0.211)	
Overestimate	0.269*** (0.075)		0.268*** (0.088)		0.307** (0.142)	
Perception gap \times Treatment		0.009** (0.004)		0.010** (0.004)		0.003 (0.008)
Perception gap		-0.012*** (0.003)		-0.012*** (0.003)		-0.010* (0.006)
Linear combination: a + b	0.020 (0.075)		-0.009 (0.084)		0.085 (0.169)	
N	1327	1327	1023	1023	304	304
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Note: The table shows OLS regression results where the dependent variable is the number of donations to the pro-black civil rights organization (the respondents were given a multiple price list where they could choose between money for themselves and \$5 to the pro-black civil rights organization in increments of \$1 from \$0 to \$5). The dependent variable has been z-scored using the mean and standard deviation in the control group. We include the following pre-specified controls in all specifications: gender, age, race (indicators for blacks and whites), region (three indicators), household size, income, education (indicator for having at least a two-year college degree), employment (indicator for having full-time work), and self-reported political affiliation (indicators for Republicans and Democrats). “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market. “Prior” is a z-scored measure of pre-treatment beliefs about the extent of racial discrimination.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.12: Treatment effects on attitudes towards pro-black policies: Pooled across experiments

	(1) Name-blind screening	(2) Black preference	(3) Black assistance	(4) Problack (Index)	(5) Name-blind screening	(6) Black preference	(7) Black assistance	(8) Problack (Index)
Treatment (a)	-0.019 (0.039)	0.002 (0.036)	-0.017 (0.038)	-0.003 (0.034)	0.003 (0.028)	0.010 (0.026)	0.007 (0.027)	0.013 (0.024)
Treatment \times Overestimate (b)	0.056 (0.053)	-0.017 (0.050)	0.033 (0.052)	0.006 (0.046)				
Overestimate	0.093** (0.038)	0.086** (0.036)	0.152*** (0.036)	0.127*** (0.033)				
Perception gap \times Treatment					-0.002 (0.002)	0.004** (0.002)	0.001 (0.002)	0.003 (0.002)
Perception gap					-0.005*** (0.001)	-0.006*** (0.001)	-0.008*** (0.001)	-0.007*** (0.001)
Linear combination: a + b	0.037 (0.036)	-0.015 (0.035)	0.016 (0.035)	0.003 (0.032)				
N	5224	5223	5220	5217	5224	5223	5220	5217
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: The table shows OLS regression results pooling observations across Experiments 1, 2 and 3. For the outcomes *Name-blind screening* (support for mandatory name-blind recruitment), *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job), and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). These outcomes are z-scored using the mean and standard deviation in the control group. *Problack (index)* is the mean of *Black preference* and *Black assistance*; this index was pre-specified. We include the following controls in all specifications: gender, age, race (indicators for blacks and whites), region (three indicators), income, education (indicator for having at least a two-year college degree), employment (indicator for having full-time work), and self-reported political affiliation (indicators for Republicans and Democrats). We also include experiment fixed effects. “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market. The variable “Perception gap” is defined as Info–Prior, i.e. 15 minus the pre-treatment estimate of the number of resumes with black-sounding names that had to be sent out to get one callback on average.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.13: Pooled results: Political heterogeneity

	(1) Name-blind screening	(2) Black preference	(3) Black assistance	(4) Problack (Index)	(5) Name-blind screening	(6) Black preference	(7) Black assistance	(8) Problack (Index)
Panel A: Non-Republicans								
Treatment	0.040 (0.047)	0.023 (0.043)	-0.004 (0.045)	0.014 (0.040)	0.031 (0.033)	-0.005 (0.031)	-0.015 (0.031)	-0.007 (0.028)
Treatment × Overestimate	-0.016 (0.062)	-0.091 (0.059)	-0.043 (0.059)	-0.073 (0.054)				
Overestimate	0.173*** (0.044)	0.149*** (0.042)	0.232*** (0.041)	0.204*** (0.038)				
Perception gap × Treatment					-0.000 (0.002)	0.004** (0.002)	0.002 (0.002)	0.004** (0.002)
Perception gap					-0.007*** (0.001)	-0.008*** (0.001)	-0.010*** (0.001)	-0.009*** (0.001)
N	3725	3724	3722	3719	3725	3724	3722	3719
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Panel B: Republicans								
Treatment	-0.141** (0.070)	-0.041 (0.065)	-0.038 (0.070)	-0.036 (0.062)	-0.054 (0.052)	0.048 (0.048)	0.066 (0.053)	0.065 (0.046)
Treatment × Overestimate (b)	0.210** (0.105)	0.189** (0.095)	0.237** (0.106)	0.220** (0.092)				
Overestimate	-0.099 (0.074)	-0.099 (0.065)	-0.070 (0.074)	-0.084 (0.064)				
Perception gap × Treatment					-0.008** (0.004)	0.001 (0.004)	-0.006 (0.004)	-0.002 (0.003)
Perception gap					0.002 (0.002)	0.000 (0.002)	-0.001 (0.003)	-0.001 (0.002)
N	1499	1499	1498	1498	1499	1499	1498	1498
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: The table shows OLS regression results pooling observations across Experiments 1, 2 and 3. Panel A shows results for Non-Republicans, while Panel B shows results for Republicans. For the outcomes *Name-blind recruitment* (support for mandatory name-blind recruitment), *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job), and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). These outcomes are z-scored using the mean and standard deviation in the control group. *Problack (index)* is the mean of *Black preference* and *Black assistance*; this index was pre-specified. We include the following pre-specified controls in all specifications: gender, age, race (indicators for blacks and whites), region (three indicators), household size, income, education (indicator for having at least a two-year college degree), employment (indicator for having full-time work), and self-reported political affiliation (indicators for Republicans and Democrats). “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market (i.e., who thought pre-treatment that resumes with black-sounding names had to be sent out more than 15 times to get one callback on average). The variable “Perception gap” is defined as Info–Prior, i.e. 15 minus the pre-treatment estimate of the number of resumes with black-sounding names that had to be sent out to get one callback on average.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.14: Results separately for each experiment with political heterogeneity

	Experiment 1 (NORC)			Experiment 2 (Research Now)			Experiment 3 (Lucid)		
	(1) Name-blind screening	(2) Black preference	(3) Black assistance	(4) Name-blind screening	(5) Black preference	(6) Black assistance	(7) Name-blind screening	(8) Black preference	(9) Black assistance
Panel A: Main specification									
Treatment (a)	0.011 (0.076)	-0.028 (0.070)	-0.015 (0.077)	-0.123* (0.064)	-0.082 (0.059)	-0.137** (0.062)	0.067 (0.064)	0.112* (0.060)	0.092 (0.061)
Treatment \times Overestimate (b)	0.079 (0.101)	-0.037 (0.094)	0.059 (0.099)	0.254*** (0.094)	0.074 (0.088)	0.140 (0.093)	-0.111 (0.085)	-0.097 (0.081)	-0.085 (0.081)
Overestimate	0.086 (0.073)	0.194*** (0.068)	0.234*** (0.071)	-0.009 (0.066)	-0.077 (0.063)	0.089 (0.065)	0.200*** (0.060)	0.179*** (0.057)	0.162*** (0.057)
N	1378	1377	1374	1720	1720	1720	2126	2126	2126
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
P-value: $a + b = 0$	0.18	0.30	0.49	0.06	0.90	0.96	0.44	0.78	0.89
Panel B: Non-Republicans									
Treatment (a)	0.112 (0.089)	-0.063 (0.085)	-0.074 (0.091)	-0.091 (0.076)	-0.015 (0.069)	-0.053 (0.071)	0.129 (0.080)	0.125* (0.074)	0.095 (0.075)
Treatment \times Overestimate (b)	-0.082 (0.115)	-0.081 (0.109)	0.030 (0.111)	0.218** (0.109)	-0.037 (0.101)	0.008 (0.104)	-0.173* (0.102)	-0.152 (0.098)	-0.148 (0.096)
Overestimate	0.201** (0.082)	0.202** (0.079)	0.265*** (0.076)	0.089 (0.076)	-0.001 (0.073)	0.202*** (0.073)	0.253*** (0.073)	0.253*** (0.070)	0.228*** (0.069)
N	1060	1059	1057	1272	1272	1272	1393	1393	1393
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
P-value: $a + b = 0$	0.68	0.04	0.51	0.10	0.48	0.56	0.50	0.67	0.38
Panel C: Republicans									
Treatment (a)	-0.181 (0.147)	0.090 (0.121)	0.117 (0.148)	-0.195* (0.118)	-0.227* (0.116)	-0.308** (0.122)	-0.043 (0.109)	0.063 (0.101)	0.080 (0.106)
Treatment \times Overestimate (b)	0.468** (0.221)	0.098 (0.185)	0.271 (0.231)	0.364* (0.191)	0.358** (0.176)	0.481** (0.202)	-0.004 (0.157)	0.068 (0.140)	0.061 (0.151)
Overestimate	-0.242 (0.164)	0.128 (0.129)	0.056 (0.178)	-0.300** (0.133)	-0.316** (0.122)	-0.289** (0.143)	0.110 (0.107)	0.011 (0.096)	0.027 (0.100)
N	318	318	317	448	448	448	733	733	733
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
P-value: $a + b = 0$	0.07	0.17	0.03	0.26	0.33	0.29	0.67	0.17	0.18

Note: The table shows OLS regression results for each experiment separately. Panel A shows results for all respondents, Panel B shows results for non-Republicans only, and Panel C show results for Republicans only. For the outcomes *Name-blind recruitment* (support for mandatory name-blind recruitment), *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job), and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). These outcomes are z-scored using the mean and standard deviation in the control group. We include the following pre-specified controls in all specifications: gender, age, race (indicators for blacks and whites), region (three indicators), household size, income, education (indicator for having at least a two-year college degree), employment (indicator for having full-time work), and self-reported political affiliation (indicators for Republicans and Democrats). “Overestimate” takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market (i.e., who thought pre-treatment that resumes with black-sounding names had to be sent out more than 15 times to get one callback on average).

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

A.4. Robustness and mechanisms

Table A.15: The association between beliefs and policy views: Robustness across elicitation techniques

	(1) Black preference	(2) Black assistance	(3) Name-blind screening	(4) Disc. against blacks
Panel A: Anchor: Black				
Beliefs about discrimination	0.208*** (0.067)	0.190*** (0.056)	0.224*** (0.041)	0.389*** (0.042)
N	347	347	347	342
Panel B: Anchor: White				
Beliefs about discrimination	0.017 (0.046)	0.082** (0.035)	0.001 (0.045)	0.157*** (0.046)
N	371	371	371	365
Panel C: Percent difference				
Beliefs about discrimination	0.126** (0.052)	0.120** (0.053)	0.040 (0.058)	0.427*** (0.043)
N	346	346	346	342

Note: This table uses data from Experiment 3. For the outcomes *Name-blind recruitment* (support for mandatory name-blind recruitment), *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job), and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). *Disc. against blacks* is the percent chance that there is racial hiring discrimination against blacks. These outcomes are z-scored using the mean and standard deviation in the control group. Panel A shows results from the prior beliefs which we anchored beliefs with the number of resumes with black-sounding names that needed to be sent out for one callback. Beliefs about discrimination is the z-scored estimate of number of resumes with white-sounding names that needed to be sent out for one callback. Panel B shows results from the prior beliefs for which we anchored beliefs with the number of resumes with white-sounding names that needed to be sent out for one callback. Beliefs about discrimination is the z-scored estimate of number of resumes with black-sounding names that needed to be sent out for one callback. Panel C shows the results for prior beliefs in which we directly measured beliefs about differences in callback rates between white-sounding names and black-sounding names.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.16: Treatment effects by prior elicitation techniques

	Anchor: Black		Anchor: White		Percent difference	
	(1) Black preference	(2) Black assistance	(3) Black preference	(4) Black assistance	(5) Black preference	(6) Black assistance
Panel A: Main effects						
Treatment	0.043 (0.077)	0.161** (0.077)	-0.013 (0.080)	-0.051 (0.076)	0.162** (0.080)	0.025 (0.081)
N	700	700	708	708	726	726
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean: Dep. var.	2.67	3.33	2.66	3.38	2.60	3.36
Panel B: Heterogeneity						
Treatment (a)	0.125 (0.098)	0.211** (0.101)	0.105 (0.154)	-0.076 (0.145)	0.134 (0.102)	0.016 (0.105)
Overestimate \times Treatment (b)	-0.212 (0.158)	-0.131 (0.155)	-0.163 (0.179)	0.035 (0.171)	0.064 (0.162)	0.021 (0.163)
Overestimate	0.357*** (0.113)	0.436*** (0.110)	0.313*** (0.114)	0.195* (0.109)	0.101 (0.121)	0.096 (0.121)
N	700	700	708	708	726	726
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Control group mean: Overestimate	0.422	0.422	0.728	0.728	0.411	0.411
P-value: a + b = 0	0.482	0.493	0.528	0.645	0.117	0.769

Note: This table uses data from Experiment 3. For the outcomes *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job), and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). These outcomes are z-scored using the mean and standard deviation in the control group. *Treatment* takes the value one for respondents who received the information treatment. *Overestimate* takes the value one for respondents who overestimate the extent of racial discrimination against blacks in the labor market. We include pre-specified controls in all regressions (the controls are listed in Table 2). Columns 1 and 2 show results from the elicitation where beliefs were anchored with the number of resumes with black-sounding names that needed to be sent out for one callback. Columns 3 and 4 show results from the elicitation where beliefs were anchored with the number of resumes with white-sounding names that needed to be sent out for one callback. Columns 5 and 6 show results for the elicitation where we directly measured beliefs about differences in callback rates between white-sounding names and black-sounding names.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

Table A.17: Experiment 2: Treatment effects – mechanism questions

	(1) Affirmative action hurts	(2) Inequality due to effort	(3) Inequality due to disc.	(4) Disc. ser. problem
Panel A: Main specification				
Treatment (a)	0.054 (0.066)	0.015 (0.061)	0.048 (0.062)	-0.022 (0.063)
Overestimate × Treatment (b)	-0.083 (0.095)	-0.121 (0.087)	-0.081 (0.090)	0.189** (0.089)
Overestimate	0.022 (0.067)	-0.080 (0.062)	0.465*** (0.064)	0.105* (0.063)
N	1720	1719	1715	1715
Controls	Yes	Yes	Yes	Yes
P-value: a + b = 0	0.669	0.087	0.607	0.008
Panel B: Political heterogeneity				
Treatment (a)	-0.076 (0.074)	-0.046 (0.071)	0.089 (0.076)	0.017 (0.076)
Overestimate × Treatment (b)	0.080 (0.104)	0.003 (0.100)	-0.099 (0.104)	0.081 (0.105)
Republican × Treatment (c)	0.441*** (0.155)	0.205 (0.139)	-0.137 (0.132)	-0.131 (0.134)
Republican × Overestimate × Treatment (d)	-0.592** (0.240)	-0.488** (0.204)	0.032 (0.211)	0.440** (0.194)
N	1720	1719	1715	1715
Controls	Yes	Yes	Yes	Yes
P-value: a + b = 0	0.953	0.543	0.891	0.182
P-value: a + c = 0	0.007	0.184	0.664	0.302
P-value: b + d = 0	0.018	0.006	0.716	0.001
P-value: a + b + c + d = 0	0.383	0.013	0.442	0.001

Note: The table shows OLS regression results where the dependent variables are indicated in each column. Responses were elicited in the second wave of Experiment 2 (the obfuscated follow-up study). *Affirmative action hurts* refers to the question of whether “affirmative action programs for the past fifty years have helped blacks” which was elicited on a scale from 1 (Strongly helped) to 7 (Strongly hurt). *Inequality due to effort* refers to the question of whether “differences in economic outcomes between whites and blacks are primarily the result of racial discrimination against blacks” which was elicited on a scale from 1 (Strongly disagree) to 7 (Strongly agree). *Inequality due to disc.* refers to the question of whether “differences in economic outcomes between whites and blacks are primarily the result of whites working harder than blacks” which was elicited on scale from 1 (Strongly disagree) to 7 (Strongly agree). *Disc. ser. problem* refers to the question of whether “racial discrimination against blacks in the labor market is a serious problem” which was elicited on a scale from 1 (Not a problem at all) to 5 (A very serious problem). All responses are z-scored using the mean and the standard deviation of the control group. Controls include gender, age, race, region, income, education, employment, political views, and confidence in prior beliefs. Only the treatment indicator and the treatment interaction terms are shown in the table.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

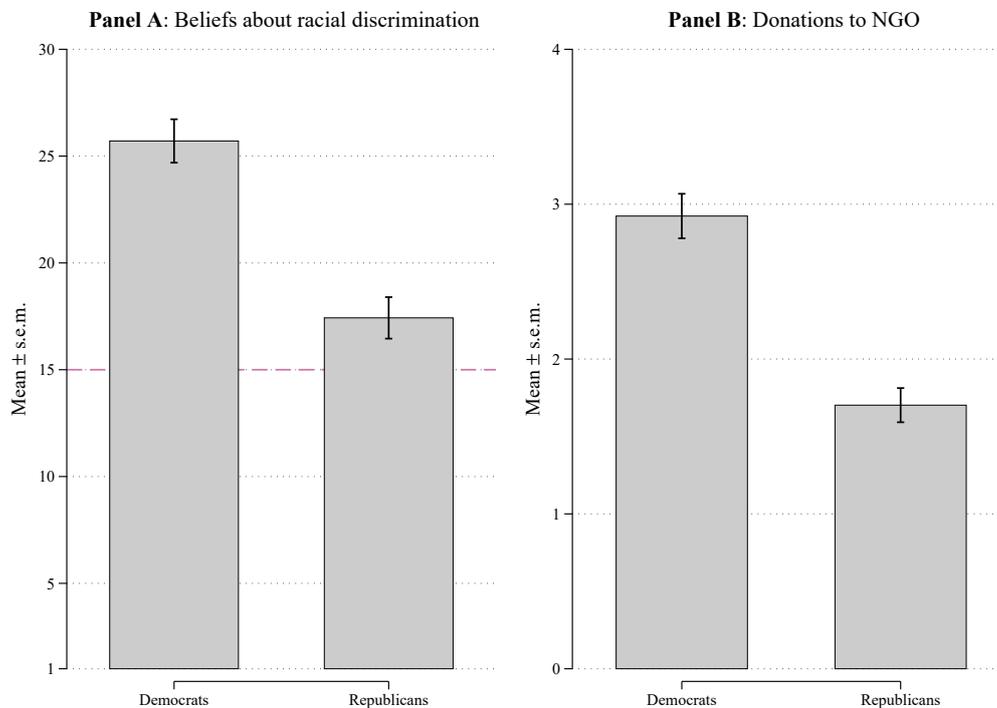
Table A.18: Correlates of willingness to pay for research evidence

	Willingness to pay	
	Raw	z-score
Republican	-0.482** (0.221)	-0.172** (0.079)
Age	0.012* (0.007)	0.004* (0.002)
Log(Income)	0.020 (0.126)	0.007 (0.045)
Black	-0.420 (0.413)	-0.150 (0.147)
White	-0.490** (0.209)	-0.175** (0.075)
College	0.309 (0.256)	0.110 (0.091)
Male	-0.459** (0.192)	-0.164** (0.069)
Prior	0.009 (0.007)	0.003 (0.002)
Confidence in prior	0.023 (0.100)	0.008 (0.036)
Mean	3.318	-0.001
Observations	861	861

Notes: This table shows OLS regressions using control group respondents from Experiment 2 (Research Now). We offered control group respondents the option to buy information about the results from the correspondence study by Bertrand and Mullainathan (2004). Willingness to pay to receive the information was elicited using a multiple price list where respondents could choose between receiving the information or varying amounts for themselves (between 10 cents and \$1). “Willingness to pay” is the number of times individuals prefer to receive information over receiving money (on a scale from 0 to 7). Column 1 shows the raw score, whereas column 2 shows the z-score (standardized using the mean and standard deviation of the responses). “Prior” is beliefs about the number of resumes with black-sounding names that had to be sent out to get one callback on average. * p<0.1, ** p<0.05, *** p<0.01. Robust standard errors in parentheses.

B. Appendix figures

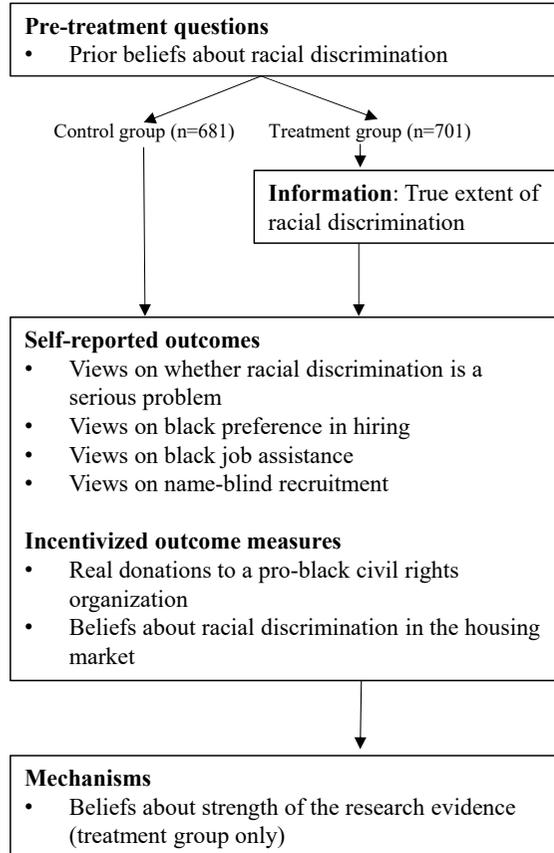
Figure A.1: Political differences in beliefs and preferences



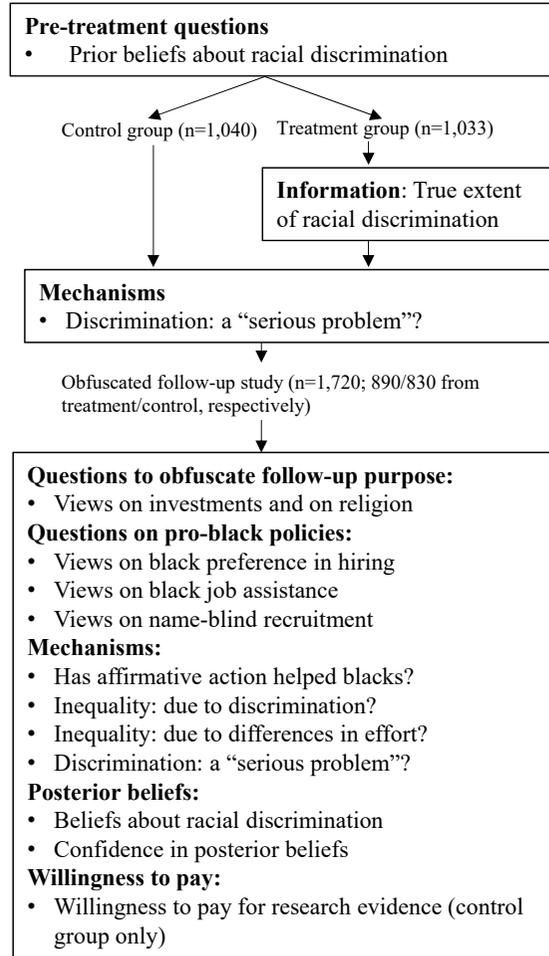
Notes: This figure uses data from Experiment 1 (the NORC sample). **Panel A** shows the mean of beliefs about how many times resumes with black-sounding names on average had to be sent out to get one callback for an interview, separately for Democrats and Republicans (the dashed line indicates the correct answer, as found in the study by Bertrand and Mullainathan, 2004). Respondents were informed that resumes with white-sounding names on average had to be sent out ten times to get one callback on average. **Panel B** shows the mean of the number of times control group respondents preferred to give \$5 to the pro-black civil rights organization over money for themselves in \$1 increments from \$0 to \$5 for Democrats and Republicans separately. Error bars indicate the standard error of the mean.

Figure A.2: Overview of experiments

Experiment 1: NORC



Experiment 2: Research Now



Experiment 3: Lucid

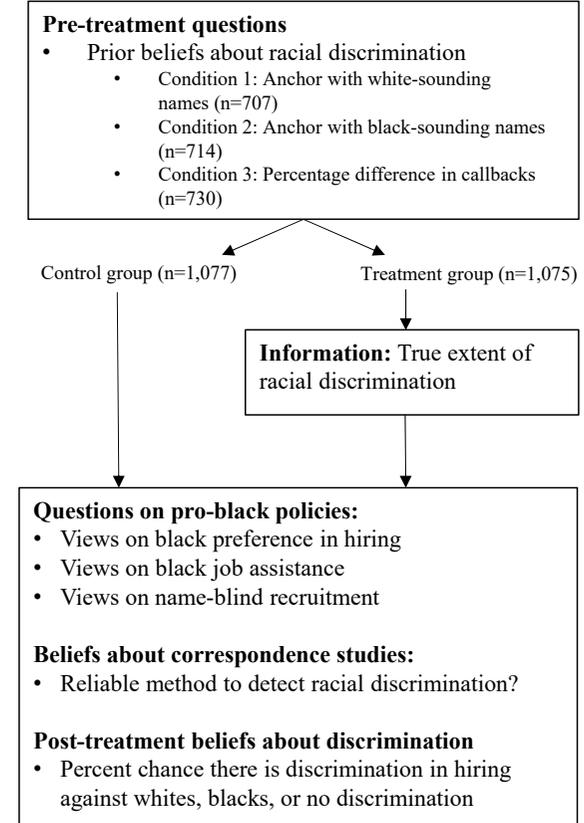
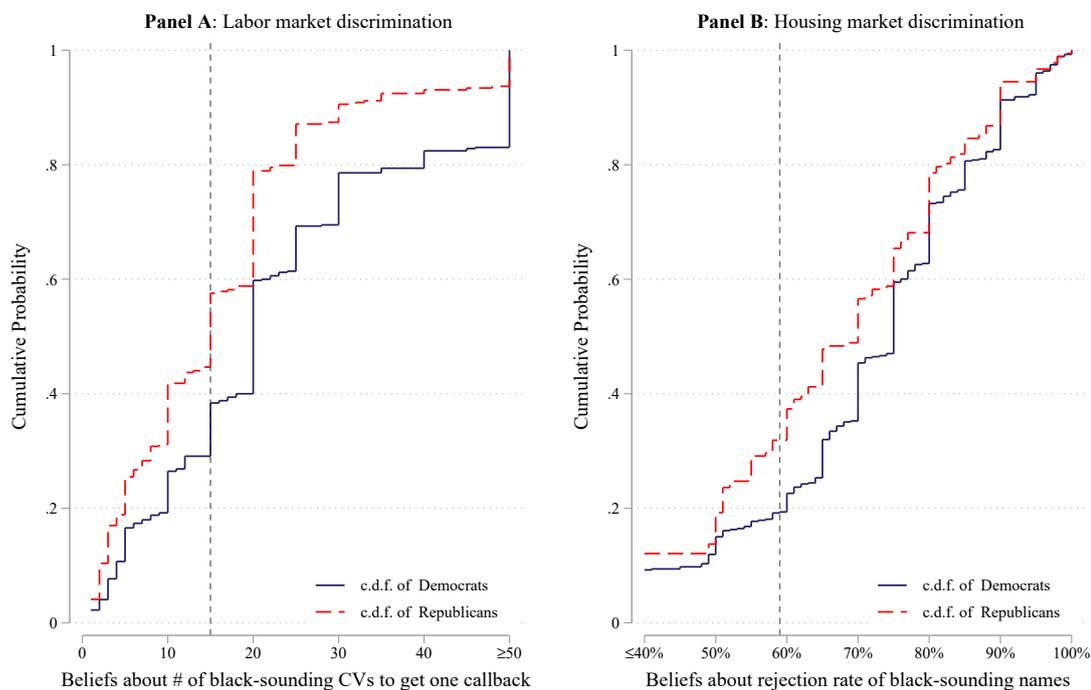
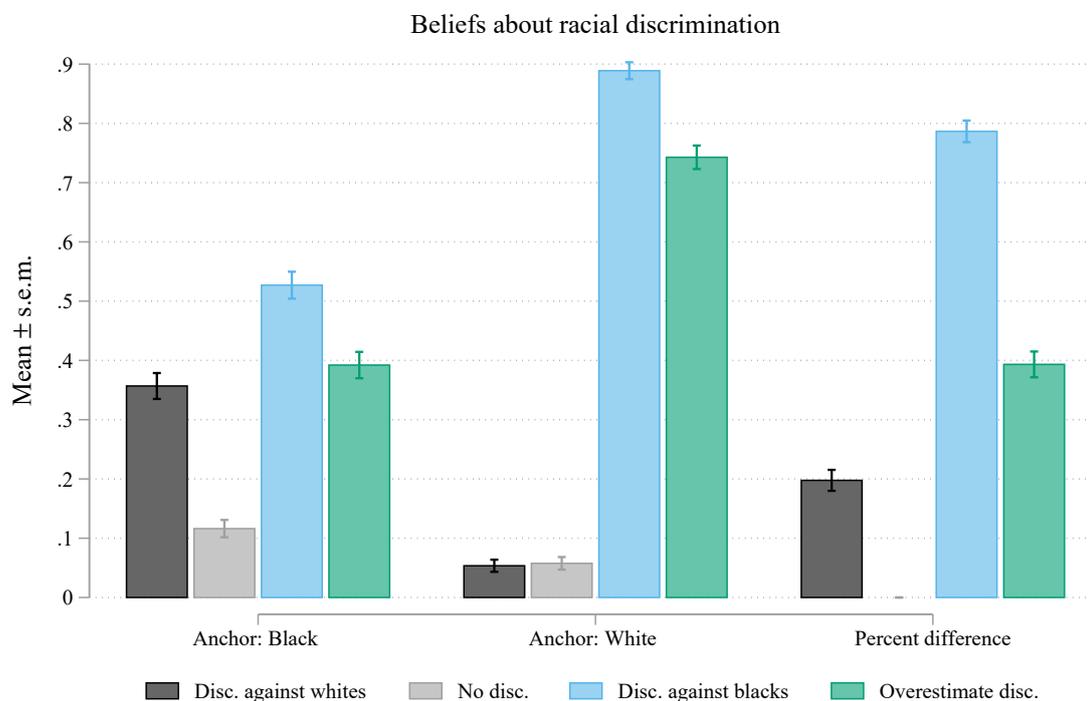


Figure A.3: Republican–Democrat differences in beliefs about racial discrimination



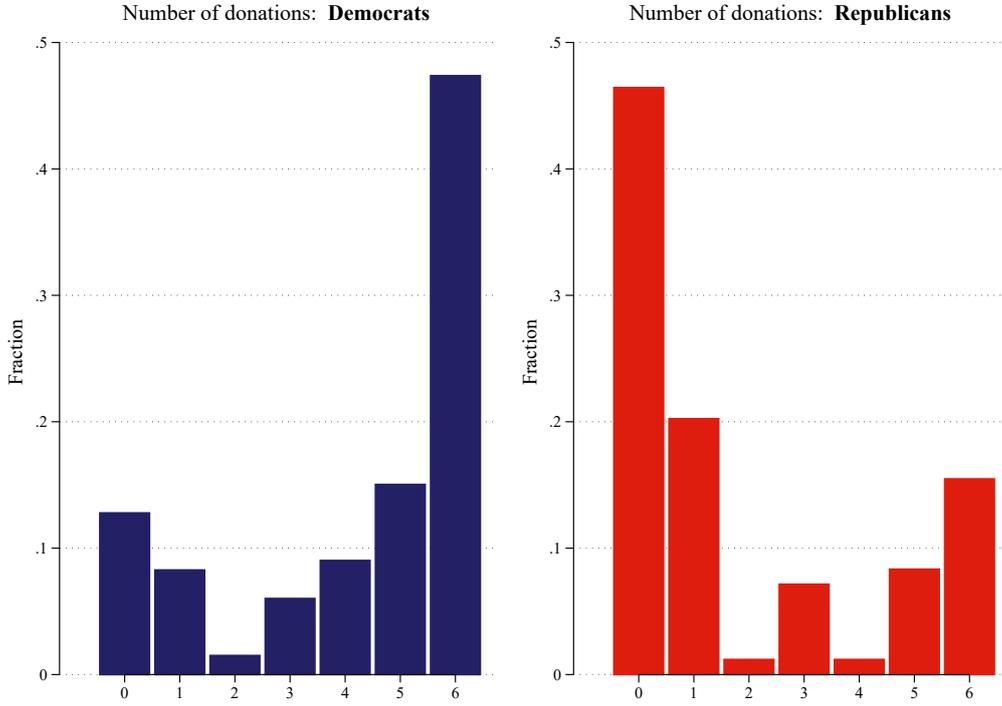
Notes: This figure uses data from Experiment 1 (the NORC sample). **Panel A** shows, separately for Republicans and Democrats, data on beliefs about how many times resumes with black-sounding names on average had to be sent out to get one callback for an interview. Respondents were informed that the corresponding number for resumes with white-sounding names was ten (as found in the study by Bertrand and Mullainathan, 2004). **Panel B** shows, separately for Republicans and Democrats, using only control group respondents, beliefs about the rejection rate on reservation requests sent from accounts with black-sounding names. Respondents were initially asked about the percent of acceptances of reservation requests for black-sounding names on Airbnb (true rate is 41 percent, as found in the study by Edelman et al., 2017). They were told that the corresponding number for white-sounding names was 49. We have recoded the values to implied rejection rates by subtracting each estimate from 100. In both panels, the dashed vertical line indicates the correct answer.

Figure A.4: Descriptives across elicitation techniques



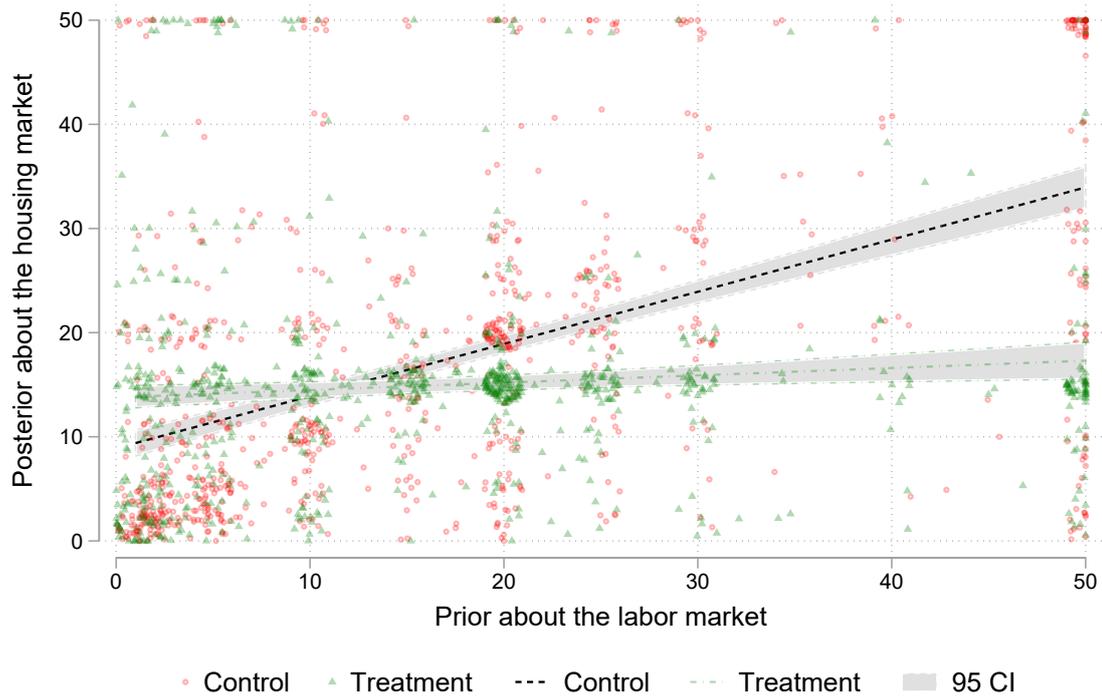
Notes: This figure shows data from Experiment 3. “Anchor black” shows results from the elicitation where beliefs were anchored with the number of resumes with black-sounding names that needed to be sent out for one callback. “Anchor white” shows results from the elicitation where beliefs were anchored with the number of resumes with white-sounding names that needed to be sent out for one callback. “Percent difference” shows results for the elicitation where we directly measured beliefs about differences in callback rates between white-sounding names and black-sounding names. The black bar indicates the fraction of respondents who thought that whites received less callbacks than blacks. The grey bar indicates the fraction of respondents who thought that blacks received equally many callbacks as whites. The blue bar indicates the fraction of respondents who thought that whites received more callbacks than blacks. The green bar indicates the fraction of respondents who thought that blacks got less callbacks compared to the findings in Bertrand and Mullainathan (2004).

Figure A.5: Republican–Democrat differences in donations behavior



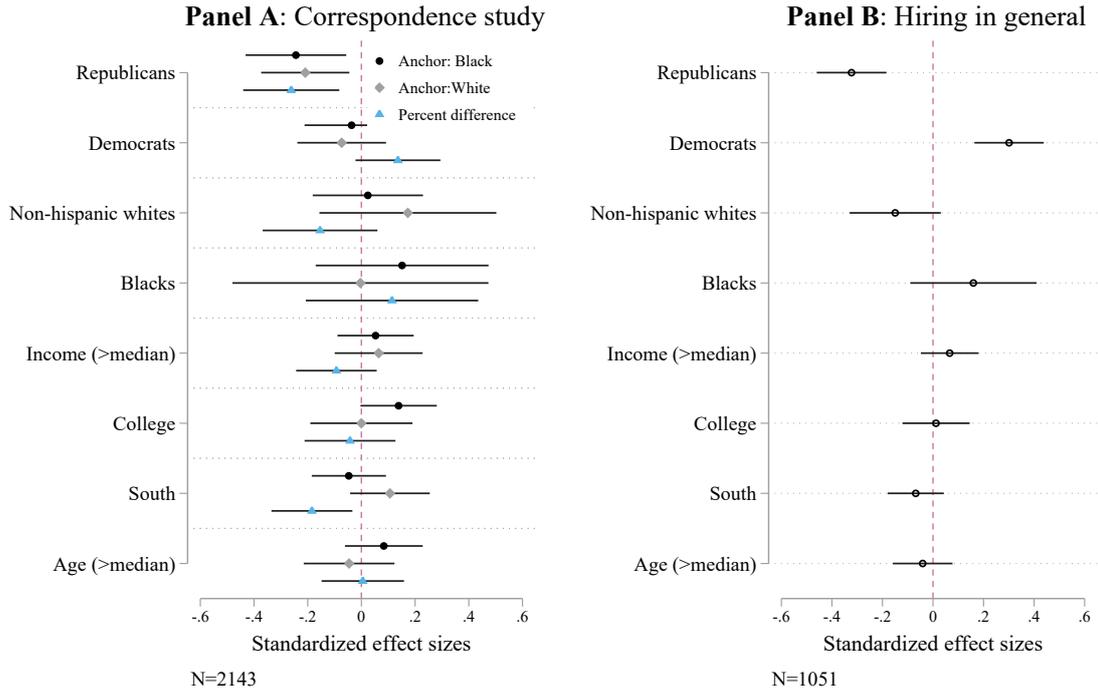
Notes: The figure, which uses data from control group respondents in Experiment 1 (NORC), shows distributions of the number of donations to the pro-black civil rights organization for self-identified Democrats and Republicans separately (the respondents were given a multiple price list where they could choose between money for themselves and \$5 to the pro-black civil rights organization in increments of \$1 from \$0 to \$5). The figure only includes respondents who completed all choices in the multiple price list.

Figure A.6: Belief updating about the labor market



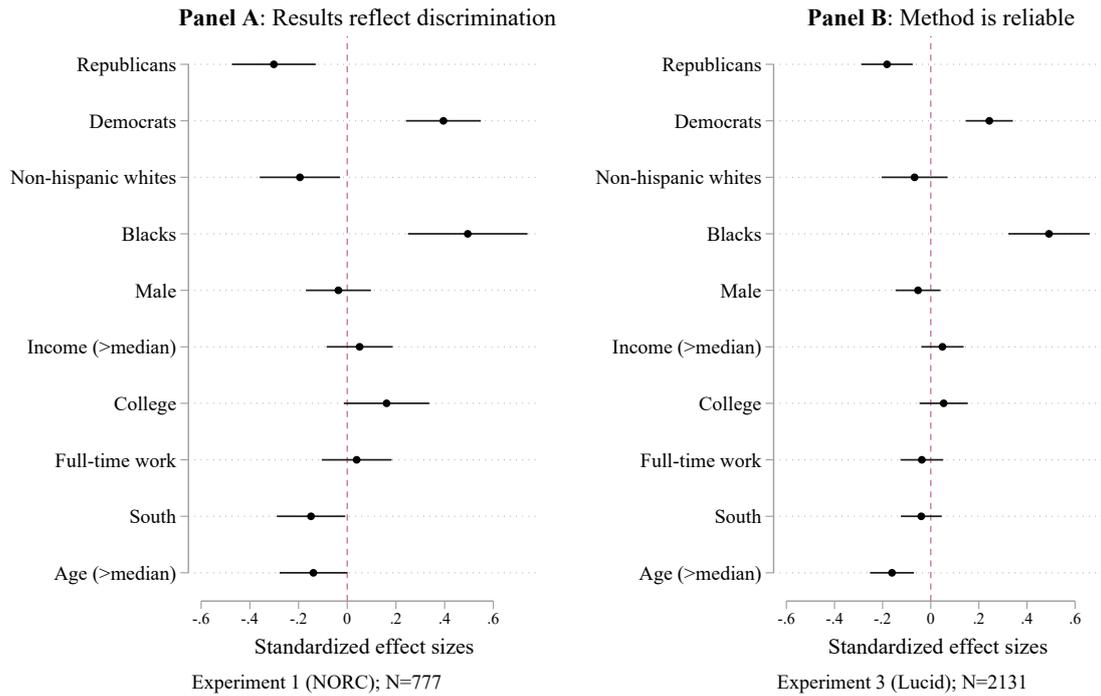
Notes: This figure, which uses data from Experiment 2 (Research Now), shows a scatterplot with prior beliefs about racial discrimination in hiring (x-axis) and posterior beliefs about racial discrimination in hiring one week later (y-axis) separately for the treatment and the control group. The prior and posterior beliefs have been top-coded at 50. Lines indicate fitted values with 95 percent confidence intervals.

Figure A.7: Correlates of beliefs about racial discrimination across elicitation techniques



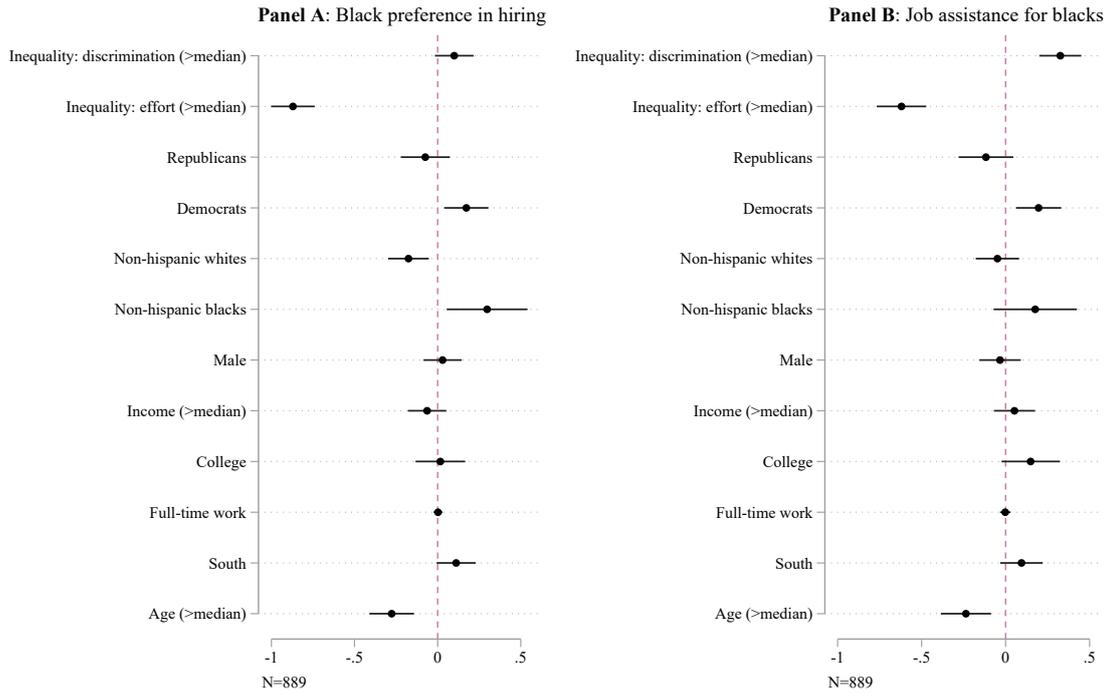
Notes: This figure uses data from Experiment 3. **Panel A** shows correlates of beliefs about callback rates in the correspondence study for the three different elicitation techniques. “Anchor black” shows results from the elicitation where beliefs were anchored with the number of resumes with black-sounding names that needed to be sent out for one callback. “Anchor white” shows results from the elicitation where beliefs were anchored with the number of resumes with white-sounding names that needed to be sent out for one callback. “Percent difference” shows results for the elicitation where we directly measured beliefs about differences in callback rates between white-sounding names and black-sounding names. **Panel B** shows correlates of more general beliefs about the percent chance that there is hiring discrimination against blacks. The dots indicate the mean values of the estimated multiple regression coefficients. Lines indicate 95 percent confidence intervals.

Figure A.8: Correlates of beliefs about correspondence studies



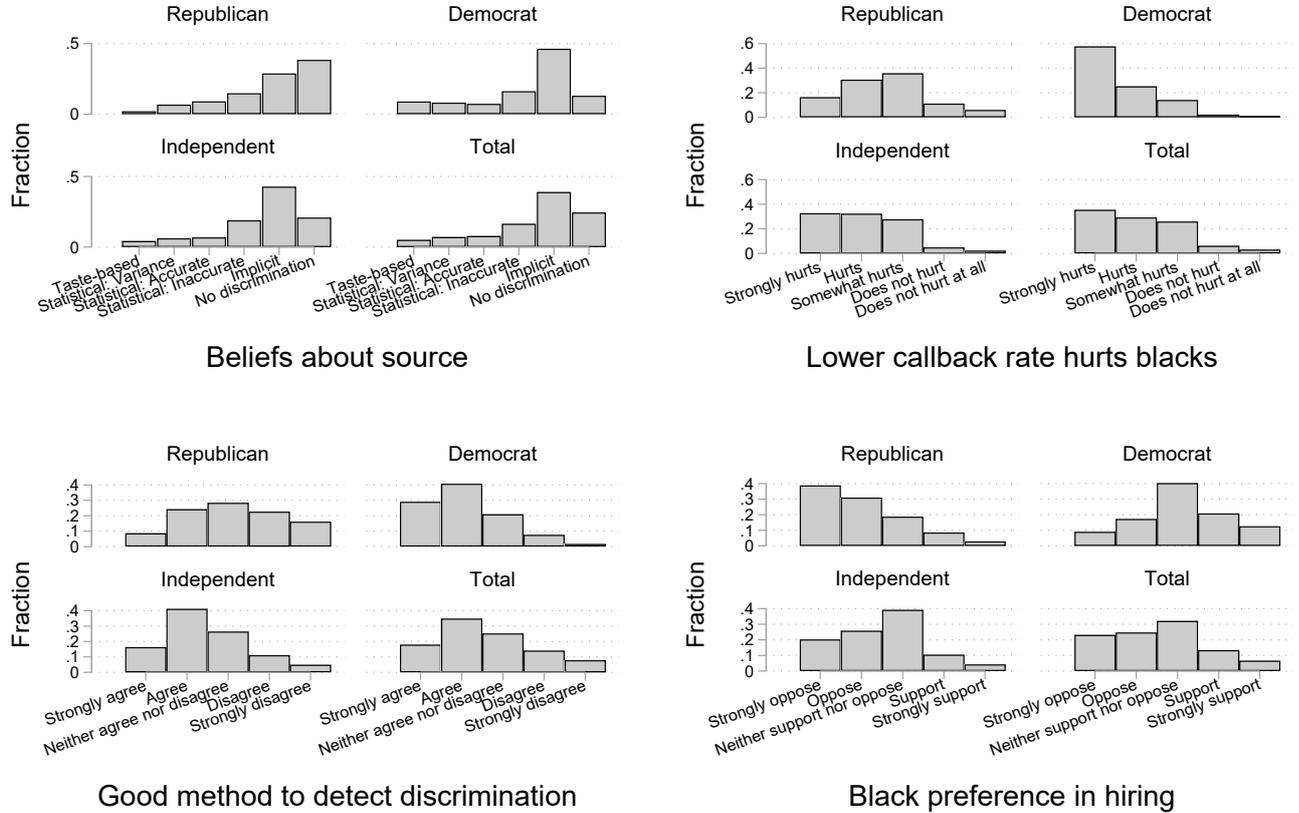
Notes: This figure uses data from Experiment 1 (NORC; Panel A) and Experiment 3 (Lucid; Panel B). The outcome variable in **Panel A** is people’s agreement with the interpretation that the findings from the correspondence study are clear evidence of discrimination against blacks in the labor market. The outcome variable in **Panel B** is people’s agreement with the following statement: “Sending out fictitious resumes to assess whether white-sounding names or black-sounding names receive more callbacks for interviews is a reliable method to detect racial discrimination in hiring.” For both questions, people respond on a 5-point scale (Strongly agree/Agree/Neither agree nor disagree/Disagree/Strongly disagree), and the outcomes are z-scored. The dots indicate the mean values of the estimated multiple regression coefficients. Lines indicate 95 percent confidence intervals.

Figure A.9: Correlates of attitudes towards pro-black policies



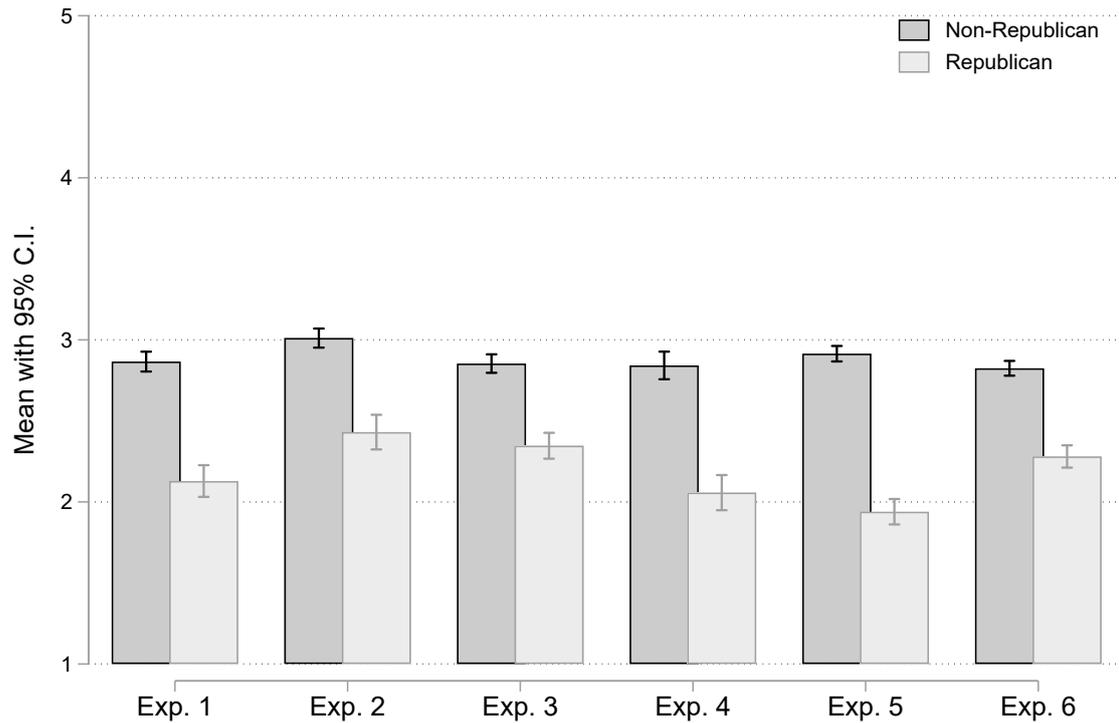
Notes: This figure uses data from Experiment 2 (Research Now). The dots indicate the mean values of the estimated multiple regression coefficients. The dependent variable in **Panel A** is support for giving black candidates preference over equally qualified white candidates in getting a job. The dependent variable in **Panel B** is support for giving qualified black candidates assistance in getting a job. Both outcomes are z-scored. “Inequality: discrimination” and “Inequality: effort” are agreements to the statements that differences in economic outcomes between blacks and whites are primarily the result of, respectively, “discrimination against blacks” and “whites working harder than blacks.” Lines indicate 95 percent confidence intervals.

Figure A.10: Results from Experiment 4



Notes: This figure uses data from Experiment 4. In **Beliefs about source**, respondents were asked “What do you think is the main reason that employers are more likely to call back applicants with white-sounding names?”. The possible answers were as follows: “They don’t want to hire black candidates because they don’t like to interact with blacks” (Taste-based) “They think the resume credentials are more informative about the skills of white job applicants than about the skills of black job applicants” (Statistical: Variance) “They correctly think that blacks on average tend to be less productive than whites” (Statistical: Accurate) “They incorrectly think that blacks on average tend to be less productive than whites” (Statistical: Inaccurate) “They subconsciously rely on negative stereotypes about blacks” (Implicit) “I don’t think it’s generally true that employers are more likely to call back applicants with white-sounding names” (No discrimination). For **Lower callback rates hurts blacks**, respondents were asked: “Do you think the lower callback rate for black-sounding names hurts blacks in the labor market?” For **Good method to detect discrimination**, respondents were asked to what extent they agreed with the statement: “Sending out fictitious resumes to assess whether white-sounding names or black-sounding names receive more callbacks for interviews is a reliable method to detect racial discrimination in hiring.” For **Black preference in hiring**, respondents were asked: “Do you support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job?”

Figure A.11: Support for black preference in hiring: Differences across experiments



Notes: This figure shows support for black preference in hiring separately for Republicans and non-Republicans for each experiment. Respondents were asked whether they “support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job.” and reported their answer on a 5-point scale ranging from 1 (Strongly oppose) to 5 (Strongly support). The bars indicate the mean value and the error bars indicate 95 percent confidence intervals. Table A.1 provides further details about each experiment.

C. Screenshots

Figure A.12: Invitation emails sent out for the experiments with Research Now

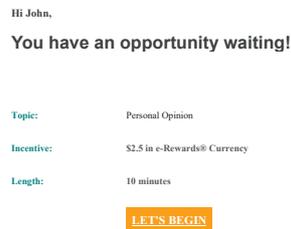


Figure A.13: Consent form in waves 1 and 2 of Experiment 2 (Research Now)

This study has received ethics clearance by the Oxford University Institutional Review Board.

If subjects have questions about this study or their rights, or if they wish to lodge a complaint or concern, they may contact us at the following email: christopher.roth@economics.ox.ac.uk

Next >>

0% 100%

Consent form

- I have read the information provided on the previous page.
- I have had the opportunity to ask questions about the study.
- I understand that I may withdraw from the study at any time.
- I understand how to raise a concern or make a complaint.
- I understand that I can only participate in this experiment once.
- I understand that close attention to the survey is required for my responses to count.

If you are 18 years of age or older, agree with the statements above, and freely consent to participate in the study, please click on the "I Agree" button to begin the experiment.

I agree I disagree

Next >>

0% 100%

This survey is conducted by a researcher from NHH Norwegian School of Economics.

In this survey, you will be asked questions on a broad range of different topics. Please pay close attention to all questions.

By continuing this survey, you acknowledge your consent to participate and that you are at least 18 years of age.

>>

D. Information about stereotypes and party views

This section discusses Experiments 5 and 6 that were discussed in Section 4.6 in the main text.

D.1. Beliefs about differences in work ethic

A centuries-old negative stereotype of blacks is the belief that they are “lazy, shiftless, and unambitious” (Gilens, 2009). One reason for why Democrats and Republicans differ in their views on pro-black policies could be that they differ in the extent to which they hold this negative stereotype.¹

In Experiment 2, we asked respondents several questions to shed light on mechanisms, including two questions on whether differences in economic outcomes between whites and blacks were primarily the result of “racial discrimination against blacks” or primarily the result of “whites working harder than blacks.” Using data from control group respondents, we show that believing that racial inequality is due to “whites working harder than blacks” is, by a large margin, the strongest predictor of attitudes towards pro-black policies (as displayed in Figure A.9). Agreeing to the statement that racial inequalities are due to “whites working harder than blacks” is associated with a 0.87 of a standard deviation lower support for black preference in hiring, conditional on controls for demographics and party affiliations ($p < 0.01$). To shed light on whether negative stereotyping of blacks causally affects attitudes towards affirmative action policies, we ran an additional experiment in which we challenge this stereotype with an information intervention.

Experimental design and sample We recruited approximately 3000 American respondents from Amazon Mechanical Turk (MTurk), an online platform commonly

¹ For a formal model of stereotypes, see Bordalo et al. (2016).

used in economic experiments (Cavallo et al., 2016; Horton et al., 2011; Kuziemko et al., 2015). We ran the experiment in October 2018 and submitted a pre-analysis plan to the same AEA RCT Registry trial as the main experiments before we started the data collection.

In the experiment, we first elicited people’s beliefs about which factors they think blacks and whites rate as least important for them in a job. We then randomized respondents into a treatment and control group. Respondents in the treatment group received information that blacks and whites both rate short working hours as the least important characteristic in a job. Respondents in the control group did not receive any information. Subsequently, we measured people’s support for pro-black policies using the same self-reported questions as in the main study.

Results In line with negative stereotyping of blacks (Gilens, 2009), the respondents think that whites are 20 percent more likely than blacks to place least weight on short working hours in a job (Table A.19). Furthermore, only 25 percent have correct beliefs that blacks actually placed the lowest weight on short working hours. But while having incorrect beliefs predicts greater opposition to pro-black policies, the information treatment does not affect support for pro-black policies. The information treatment also does not shift beliefs about whether differences in economic outcomes between blacks and whites are “primarily the result of whites working harder than blacks,” suggesting that the treatment is ineffective in challenging the stereotype of “lazy blacks.” Given our large sample size, we take this as suggestive evidence that beliefs governing racial stereotypes are much less responsive to new information than beliefs about racial discrimination. Furthermore, this result emphasizes that views on pro-black policies are generally very unresponsive to new information.

Table A.19: Experiment 5: Treatment effects of information about racial stereotypes

	(1) Black preference	(2) Black assistance	(3) Problack (Index)	(4) Inequality: effort
Panel A: Main specification				
Treatment	-0.001 (0.032)	0.012 (0.033)	0.006 (0.028)	0.040 (0.032)
Panel B: Heterogeneity				
Treatment (a)	0.02 (0.04)	0.03 (0.04)	0.02 (0.03)	0.04 (0.04)
Prior × Treatment (b)	-0.09 (0.07)	-0.06 (0.07)	-0.07 (0.06)	0.01 (0.07)
Prior	0.18*** (0.05)	0.11** (0.05)	0.15*** (0.05)	-0.13*** (0.05)
N	2999	2999	2999	2999
Controls	Yes	Yes	Yes	Yes
P-value: a + b = 0	0.29	0.61	0.37	0.43

Note: The table shows OLS regression results from Experiment 5 (MTurk). The dependent variables are indicated in each column. For the outcomes *Black preference* (support for giving qualified black candidates preference over equally qualified white candidates in getting a job) and *Black assistance* (support for giving qualified black candidates assistance in getting a job), answers were given on a scale from 1: “Strongly oppose” to 5: “Strongly support.” These outcomes are z-scored using the mean and standard deviation in the control group. *Problack (index)* is the mean of *Black preference* and *Black assistance*; this index was pre-specified. For the outcome “Inequality: effort” (agreement to the statement that differences in economic outcomes between blacks and whites are due to whites working harder than blacks), answers were given on a scale from 1 (Strongly disagree) to 7 (Strongly agree) and then z-scored. *Prior* is indicator taking the value one for respondents who thought that blacks were most likely to rank “Working hours are short, lots of free time” as the least important characteristic in a job. Controls were pre-specified and include the prior, two racial indicators (black and white), a gender indicator, a college indicator, age, log income, and two indicators for political status (Democrats and Republicans).

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

D.2. Information about party views on affirmative action

During the last four decades, political polarization in beliefs about whether differences in economic outcomes between blacks and whites are “mainly due to discrimination” has strongly increased (Smith et al., 2014). This shift in beliefs is part of a broader trend in which American politics has become more polarized along partisan lines than at any point in recent history.² This experiment sheds light on the role of political identity (Bursztyn et al., 2019) and the importance of party cues (Brader and Tucker, 2012) in shaping policy views by providing information about how the Republican and Democratic parties differ in their support for affirmative action policies.

Experimental sample and design We recruited 4000 respondents in collaboration with Research Now, the same market research company as used in Experiment 2. The sample was constructed to be representative of the US population in terms of age, sex, and region. We ran the experiment in July 2018, and we submitted a pre-analysis plan to the same AEA RCT Registry trial as the main experiments before we started the data collection.³

We randomly assigned respondents into a control group and a treatment group. For respondents in the treatment group, we added the following introductory sentence to the question on whether they support affirmative action in hiring: “In contrast to the Democratic Party, the Republican Party generally opposes all forms of special treatment based on race.” In the main specification, we focused on the 2,737 respondents who self-identify as either Democrats or Republicans. We

² <http://www.people-press.org/2014/06/12/political-polarization-in-the-american-public>, accessed September 28, 2020.

³ Instructions are provided in Section F.4 of the Online Appendix.

hypothesized in the pre-analysis plan that this treatment would polarize attitudes by making Democrats more supportive of pro-black policies while making Republicans less supportive.

Results The treatment has essentially no impact on attitudes for either Democrats or Republicans (Table A.20). Given our large sample size and the importance of political identity and party cues documented in previous research (Brader and Tucker, 2012; Cappelen et al., 2017), this finding again underscores the point that views on pro-black policies are hard to move with information.

Table A.20: Experiment 6: Treatment effects of information about party views

	(1)	(2)	(3)	(4)
Treatment	0.04 (0.05)	0.03 (0.05)	-0.04 (0.05)	-0.03 (0.05)
Republicans	-0.61*** (0.05)	-0.61*** (0.05)	-0.20*** (0.05)	-0.17*** (0.05)
Treatment × Republicans	-0.06 (0.07)	-0.03 (0.07)	0.02 (0.07)	0.03 (0.07)
Democrats			0.41*** (0.05)	0.44*** (0.05)
Treatment × Democrats			0.08 (0.07)	0.06 (0.07)
N	2737	2737	4000	4000
Controls	No	Yes	No	Yes

Note: The table shows OLS regressions from Experiment 6 (Research Now). The dependent variable is support for “government and private programs that give qualified black and other racial minority candidates preference over equally qualified white candidates in getting a job.” Answers were given on a scale from 1 (Strongly oppose) to 5 (Strongly support). We have z-scored the responses by the mean and standard deviation in the control group. In the treatment group, we informed respondents about party views on affirmative action as follows: “In contrast to the Democratic Party, the Republican Party generally opposes all forms of special treatment based on race.” In even-numbered columns, we include the following pre-specified controls: gender, age, and education. In line with the pre-analysis, we exclude Independents from the regression in columns 1–2 as the treatment was tailored to affect attitudes for Republicans and Democrats. In columns 3–4, we add interaction terms between the treatment and Democrats and add Independents to the regressions. The sample was recruited from Research Now and is representative of the US population on the following observable characteristics: age, sex, and region.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses.

E. Beliefs about the consequences of callback differences

One potential reason for why respondents do not change their views on pro-black policies in response to the research evidence could be that they think that lower callback rates do not hurt blacks in the labor market. For instance, if discrimination is due to prejudiced employers, people might think that racial segregation of labor markets allows blacks to avoid contact with prejudiced employers. Correspondence studies have therefore been criticized for surveying the average firm rather than the marginal discriminatory firm (Heckman, 1998). People might also think that blacks easily can compensate for lower callback rates by sending out more resumes.

Using data from Experiment 4, we shed light on this issue by asking our respondents whether they think that the lower callback rate for black-sounding names hurts blacks in the labor market. We find that 91 percent of respondents respond either that the differences in callback rates strongly hurt (35 percent), hurt (29 percent) or somewhat hurt (26 percent) blacks in the labor market. Even 83 percent of Republicans think that the differences in callback rates strongly hurt (16 percent), hurt (31 percent) or somewhat hurt (36 percent) blacks in the labor market (Figure A.10). Overall, these results demonstrate that it is very unlikely that respondents do not respond to the information because they think that the callback differential does not hurt blacks in the labor market.

F. Instructions

F.1. Experiment 1 (NORC)

F.1.1. Elicitation of beliefs about racial discrimination

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did so by sending out fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names like “Tanisha” and “Kareem”. The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer whether they were white or black.

Resumes with white-sounding names had to be sent out on average 10 times to get one callback for an interview.

What do you think?

How many times do you think resumes with black-sounding names on average had to be sent out to get one callback for an interview?

I think resumes with black-sounding names on average had to be sent out

times to get one callback for an interview.

If your answer is the same as what the researchers found, you will be rewarded a bonus of \$2 (2,000 AmeriPoints) in addition to your current incentive of 2,000 AmeriPoints.

F.1.2. Treatment screen

The researchers found that resumes with black-sounding names on average had to be sent out 15 times to get one callback for an interview.

Since resumes with white-sounding names on average only had to be sent out 10 times to get one callback for an interview, this means that employers were 50 percent more likely to give callbacks to applicants with white-sounding names compared to applicants with black-sounding names.

F.1.3. Self-reported outcomes

In the United States today, do you think that racial discrimination against blacks in the labor market is a serious problem? [A very serious problem/A problem/A problem/A small problem/Not a problem at all]

Do you support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

Do you support or oppose government and private programs that give qualified black candidates assistance in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

Name-blind recruitment has been suggested as a way to reduce racial discrimination in the labor market by hiding the names of the job applicants from their resumes. Do you support or oppose mandatory name-blind recruitment for hiring in public and private jobs? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

F.1.4. Behavioral measure: Donation

In Washington, D.C., several civil rights organizations work to protect individuals from discrimination in society. One of these organizations, the *Lawyers' Committee*

for Civil Rights, tries to help African Americans. One of the organization's key initiatives aims to reduce racial discrimination in the workplace by lobbying for political reforms.

Below, you are given the opportunity to financially support the *Lawyers' Committee for Civil Rights*.

Your decision

For each of the 6 choices below, you decide whether the *Lawyers' Committee for Civil Rights* should get money or whether *you* should get money (\$1 equals 1000 AmeriPoints).

We will randomly implement your decision for *one* of these choices, which involve real money, so please consider each choice carefully. Each decision has the same chance of being implemented.

- | | | | |
|--------------------------|-----------------------|-----------------------|------------|
| \$5 for the organization | <input type="radio"/> | <input type="radio"/> | \$0 for me |
| \$5 for the organization | <input type="radio"/> | <input type="radio"/> | \$1 for me |
| \$5 for the organization | <input type="radio"/> | <input type="radio"/> | \$2 for me |
| \$5 for the organization | <input type="radio"/> | <input type="radio"/> | \$3 for me |
| \$5 for the organization | <input type="radio"/> | <input type="radio"/> | \$4 for me |
| \$5 for the organization | <input type="radio"/> | <input type="radio"/> | \$5 for me |

Note: NORC is a non-partisan research organization and has no association with the Lawyers' Committee for Civil Rights. NORC and the AmeriSpeak Panel do not endorse political or charitable causes.

F.1.5. Belief extrapolation: Discrimination in the housing market

Researchers from Harvard Business School conducted an experiment to study racial discrimination in the rental market by sending out reservation requests from invented accounts to hosts on Airbnb, a website for private rental accommodations. The requests were exactly the same except for one thing: the name of the person who sent the request. Half of the requests came from typically white-sounding names, while the other half came from typically black-sounding names. The idea was that the hosts would use the applicants' name to infer whether the reservation requests came from white or black requesters.

The researchers found that reservation requests from white-sounding names were accepted 49 percent of the time.

What do you think?

How many percent of the time do you think reservation requests from black-sounding names were accepted?

I think reservation requests from black-sounding names were accepted percent of the time.

If your answer is within 2 percentage points of what the researchers found, you will be rewarded a bonus of \$2 (2,000 AmeriPoints) in addition to your current incentive of 2,000 AmeriPoints.

F.1.6. Beliefs about strength of the evidence: Treatment group only

The researchers behind the study on labor market discrimination described earlier in this survey interpreted their findings as clear evidence of discrimination against blacks in the labor market.

To what extent do you agree or disagree with this interpretation of their findings?
[Strongly agree/Agree/Neither agree nor disagree/Disagree/Strongly disagree]

F.2. Instructions: Experiment 2 – first wave (Research Now)

F.2.1. Consent Form

This study has received ethics clearance by the Oxford University Institutional Review Board.

If subjects have questions about this study or their rights, or if they wish to lodge a complaint or concern, they may contact us at the following email:

christopher.roth@economics.ox.ac.uk.

{page break}

Consent form

I have read the information provided on the previous page.

I understand that I may withdraw from the study at any time.

I have had the opportunity to ask questions about the study.

I understand how to raise a concern or make a complaint.

I understand that I can only participate in this experiment once.

I understand that close attention to the survey is required for my responses to count.

If you are 18 years of age or older, agree with the statements above, and freely consent to participate in the study, please click on the “I agree” button to begin the experiment.

I agree

I disagree

F.2.2. Elicitation of beliefs about racial discrimination

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did so by sending out fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names like “Tanisha” and “Kareem”.

The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer whether they were white or black.

Resumes with **white-sounding** names had to be sent out on average **10 times** to get one callback for an interview.

What do you think?

How many times do you think resumes with **black-sounding** names on average had to be sent out to get one callback for an interview?

I think resumes with black-sounding names on average had to be sent out times to get one callback for an interview.

If your answer is the same as what the researchers found, you will be rewarded a **bonus of \$2** in panel currency.

F.2.3. Confidence in priors

How sure are you about your answer to the previous question?

Very sure

Sure

Somewhat sure

Unsure

Very unsure

F.2.4. Treatment screen

The researchers found that resumes with black-sounding names on average had to be sent out **15 times** to get one callback for an interview.

Since resumes with white-sounding names on average only had to be sent out 10 times to get one callback for an interview, this means that employers were **50 percent** more likely to give callbacks to applicants with white-sounding names than applicants with black-sounding names.

F.2.5. Perceptions of discrimination

In the United States today, do you think that racial discrimination against blacks in the labor market is a serious problem? [A very serious problem/A problem/A problem/A small problem/Not a problem at all]

F.3. Instructions: Experiment 2 – second wave (Research Now)

F.3.1. Introduction

This survey is conducted by a researcher from NHH Norwegian School of Economics. In this survey, you will be asked questions on a broad range of different topics. Please pay close attention to all questions.

By continuing this survey, you acknowledge your consent to participate and that you are at least 18 years of age.

F.3.2. Obfuscation: Views on investments

Which of the following do you think is the best long-term investment: bonds, real estate, saving accounts, stock or mutual funds, or gold? [Bonds, Real estate, Saving accounts, Stock or mutual funds, Gold]

{page break}

Do you, personally, or jointly with a spouse, have any money invested in the stock market right now – either in an individual stock, a stock mutual fund, or in a self-directed 401-K or IRA? [Yes, No, Do not know]

F.3.3. Obfuscation: Views on religion

How important would you say religion is in your own life – very important, fairly important, or not very important? [Very important, Fairly important, Not very important]

{page break}

At the present time, do you think religion as a whole is increasing its influence on American life or losing its influence? [Increasing/Decreasing/No opinion]

F.3.4. Self-reported outcomes

Do you support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

{page break}

Do you support or oppose government and private programs that give qualified black candidates assistance in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

{page break}

Name-blind recruitment has been suggested as a way to reduce racial discrimination in the labor market by hiding the names of the job applicants from their resumes. Do you support or oppose mandatory name-blind recruitment for hiring in public and private jobs? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

F.3.5. Mechanisms

Overall, do you think affirmative action programs for the past fifty years have helped blacks, hurt them, or had no effect one way or the other? [Strongly helped/Helped/Somewhat helped/Neither helped nor hurt/Somewhat hurt/Hurt/Strongly hurt]

To what extent do you agree with the following statement: “Differences in economic outcomes between whites and blacks are primarily the result of racial discrimination against blacks.” [Strongly agree/Agree/Somewhat agree/Neither agree nor disagree/Somewhat disagree/Disagree/Strongly disagree]

To what extent do you agree with the following statement: “Differences in economic outcomes between whites and blacks are primarily the result of whites working harder than blacks.” [Strongly agree/Agree/Somewhat agree/Neither agree nor disagree/Somewhat disagree/Disagree/Strongly disagree]

{page break}

In the United States today, do you think that racial discrimination against blacks in the labor market is a serious problem? [A very serious problem/A problem/A problem/A small problem/Not a problem at all]

F.3.6. Elicitation of posterior about labor market discrimination

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did

so by sending out fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names like “Tanisha” and “Kareem”.

The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer whether they were white or black.

Resumes with **white-sounding** names had to be sent out on average **10 times** to get one callback for an interview.

What do you think?

How many times do you think resumes with **black-sounding** names on average had to be sent out to get one callback for an interview?

I think resumes with black-sounding names on average had to be sent out times to get one callback for an interview.

If your answer is the same as what the researchers found, you will be rewarded a **bonus of \$2** in panel currency.

F.3.7. Confidence in posteriors

How sure are you about your answer to the previous question? [Very sure/Sure/Somewhat sure/Unsure/Very unsure]

F.3.8. Willingness to pay for the information (control group only)

We just explained to you the details of a study which tested for racial discrimination in the labor market.

For each of the seven choices below, you decide whether you would like to receive

more information about the results from the study or whether you would like to receive money.

If you decide to receive the information about the results of the study, we will provide you with a short summary of the results, including information on the number of times resumes with black-sounding names had to be sent out in order to get one callback. If you decide to receive the information about the results of the study, we will also provide you with a link to the research study which further describes the methodology, implementation of the experiment, and discusses the research results.

We will randomly implement your decision for *one* of these choices after the study has ended, so please consider each choice carefully. Each decision has the same chance of being implemented.

- | | | | |
|-------------|-----------------------|-----------------------|---------------|
| Information | <input type="radio"/> | <input type="radio"/> | \$0.10 for me |
| Information | <input type="radio"/> | <input type="radio"/> | \$0.20 for me |
| Information | <input type="radio"/> | <input type="radio"/> | \$0.30 for me |
| Information | <input type="radio"/> | <input type="radio"/> | \$0.40 for me |
| Information | <input type="radio"/> | <input type="radio"/> | \$0.50 for me |
| Information | <input type="radio"/> | <input type="radio"/> | \$0.75 for me |
| Information | <input type="radio"/> | <input type="radio"/> | \$1 for me |

F.3.9. Information provision (depending on people's choices)

The researchers found that resumes with black-sounding names on average had to be sent out 15 times to get one callback for an interview.

Since resumes with white-sounding names on average only had to be sent out 10 times to get one callback for an interview, this means that employers were 50 percent more likely to give callbacks to applicants with white-sounding names compared to applicants with black-sounding names.

http://www2.econ.iastate.edu/classes/econ321/orazem/bertrand_emily.pdf

F.4. Experiment 3 (Lucid)

F.4.1. Elicitation: Anchor black

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did so by sending out over 4800 fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names like “Tanisha” and “Kareem”.

The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer whether they were white or black.

Resumes with **black-sounding** names had to be sent out on average 15 times to get one callback for an interview.

What do you think?

How many times do you think resumes with **white-sounding names** on average had to be sent out to get one callback for an interview?

I think that a resume with a white-sounding name on average had to be sent out times to get a callback for an interview.

F.4.2. Elicitation: Anchor white

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did so by sending out over 4800 fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names like “Tanisha” and “Kareem”.

The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer whether they were white or black.

Resumes with **white-sounding** names had to be sent out on average **10 times** to get one callback for an interview.

What do you think?

How many times do you think resumes with **black-sounding names** on average had to be sent out to get one callback for an interview?

I think that a resume with a black-sounding name on average had to be sent out times to get a callback for an interview.

F.4.3. Elicitation: Percentage difference

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did so by sending out over 4800 fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names like “Tanisha” and “Kareem”.

The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer

whether they were white or black.

What do you think that the study found?

I think that white-sounding names received more callbacks for interviews than black-sounding names.

I think that black-sounding names received more callbacks for interviews than white-sounding names.

F.4.4. Percentage differences: Screen for respondents who thought white-sounding names were more likely to get a callback

You said that you think that white-sounding names received more callbacks for interviews than black-sounding names. How many percent more callbacks for interviews do you think white-sounding names received?

I think white-sounding names received percent more callbacks.

F.4.5. Percentage differences: Screen for respondents who thought black-sounding names were more likely to get a callback

You said that you think that black-sounding names received more callbacks for interviews than white-sounding names. How many percent more callbacks for interviews do you think black-sounding names received?

I think black-sounding names received percent more callbacks.

F.4.6. Information treatment

For your information, the study found that **white-sounding names received 50 percent more callbacks** for interviews than black-sounding names.

F.4.7. Views on pro-black policies

Do you support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job?

[Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

Do you support or oppose government and private programs that give qualified black candidates assistance in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

Name-blind recruitment has been suggested as a way to reduce racial discrimination in the labor market by hiding the names of the job applicants from their resumes. Do you support or oppose mandatory name-blind recruitment for hiring in public and private jobs? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

F.4.8. Beliefs about correspondence studies

To what extent do you agree with the following statement: “Sending out fictitious resumes to assess whether white-sounding names or black-sounding names receive more callbacks for interviews is a reliable method to detect racial discrimination in hiring.” [Strongly agree/Agree/Neither agree nor disagree/Disagree/Strongly disagree]

F.4.9. Post-treatment beliefs about discrimination

What would you say is the likelihood that...

(Please note: The numbers need to add up to 100%)

there is discrimination against blacks in hiring. %

there is discrimination against whites in hiring. %

there is no racial discrimination in hiring. %

F.5. Instructions: Experiment 4: Interpretation of correspondence studies

F.5.1. Beliefs about sources of discrimination

Researchers from Harvard University and the University of Chicago conducted an experiment to study racial discrimination in the labor market. They did so by sending out over 4800 fictitious resumes to help-wanted ads in Boston and Chicago newspapers.

The resumes were exactly the same except for one thing: the name of the job applicant. Half of the resumes had typically white-sounding names like “Carrie” and “Todd”. The other half of the resumes had typically black-sounding names like “Tanisha” and “Kareem”.

The idea was to make sure that the applicants were seen as having identical qualifications, but that the employers would use the applicants’ names to infer whether they were white or black.

The researchers found that white-sounding names received 50 percent more callbacks for interviews than black-sounding names.

What do you think is the main reason that employers are more likely to call back applicants with white-sounding names?

They don’t want to hire black candidates because they don’t like to interact with blacks

They think the resume credentials are more informative about the skills of white job applicants than about the skills of black job applicants

They correctly think that blacks on average tend to be less productive than whites

They incorrectly think that blacks on average tend to be less productive than whites

They subconsciously rely on negative stereotypes about blacks

I don't think it's generally true that employers are more likely to call back applicants with white-sounding names

F.5.2. Beliefs about whether correspondence studies is a good method

To what extent do you agree with the following statement: "Sending out fictitious resumes to assess whether white-sounding names or black-sounding names receive more callbacks for interviews is a reliable method to detect racial discrimination in hiring." [Strongly agree/Agree/Neither agree nor disagree/Disagree/Strongly disagree]

F.5.3. Beliefs about whether discrimination hurts blacks

Employers are on average 50 percent more likely to call back applications with white-sounding names than applicants with black-sounding names. Do you think the lower callback rate for black-sounding names hurts blacks in the labor market? [Strongly hurts/Hurts/Somewhat hurts/Does not hurt/Does not hurt at all]

F.5.4. Policy preferences

Do you support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

F.6. Instructions: Experiment 5: Racial stereotypes

F.6.1. Pre-treatment beliefs

In this survey, we will ask you some questions about whites and blacks in America.

Throughout this survey, we will refer to non-Hispanic whites and non-Hispanic blacks as whites and blacks, respectively.

{page break}



The General Social Survey (GSS) is a large and representative survey of Americans. In the survey, people were asked to rank the importance of the following five job characteristics (from least important to most important):

- High income
- No danger of being fired
- Working hours are short, lots of free time
- Chances for advancement
- Work that is important and gives a feeling of accomplishment

Among **whites**, which response do you think was most commonly chosen as the **least** important characteristic of a job?

High income

No danger of being fired

Working hours are short, lots of free time

Chances for advancement

Work that is important and gives a feeling of accomplishment

Among **blacks**, which response do you think was most commonly chosen as the **least** important characteristic of a job?

High income

No danger of being fired

Working hours are short, lots of free time

Chances for advancement

Work that is important and gives a feeling of accomplishment

F.6.2. Information treatment



The actual results on which response people most commonly chose as **least** important characteristic of a job were as follows:

Among **whites**, the response “Working hours are short, lots of free time” was most commonly chosen as the **least** important characteristic of a job.

Among **blacks**, the response “Working hours are short, lots of free time” was most commonly chosen as the **least** important characteristic of a job.

Source: The General Social Survey

F.6.3. Views on pro-black policies

We will now ask you a few questions about your attitudes towards policies to help blacks in the labor market.

{page break}

Do you support or oppose government and private programs that give qualified black candidates preference over equally qualified white candidates in getting a job?

[Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

{page break, note: We randomize the order of these two questions}

Do you support or oppose government and private programs that give qualified black candidates assistance in getting a job?

[Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

F.6.4. Post-treatment beliefs

To what extent do you agree with the following statement:

“Differences in economic outcomes between whites and blacks are primarily the result of whites working harder than blacks.” [Strongly agree/Agree/Somewhat agree/Neither agree nor disagree/Somewhat disagree/Disagree/Strongly disagree]

F.7. Instructions: Experiment 6: Political identity and party cues

F.7.1. Treatment group

A much debated issue is whether blacks and other racial minorities should get preference over equally qualified white candidates in getting a job. In contrast to the Democratic Party, the Republican Party generally opposes all forms of special treatment based on race. We are interested in what you think about this issue.

Do you support or oppose government and private programs that give qualified black and other racial minority candidates preference over equally qualified white candidates in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]

F.7.2. Control group group

A much debated issue is whether blacks and other racial minorities should get preference over equally qualified white candidates in getting a job. We are interested in what you think about this issue.

Do you support or oppose government and private programs that give qualified black and other racial minority candidates preference over equally qualified white candidates in getting a job?

F.7.3. Outcome measure

Do you support or oppose government and private programs that give qualified black and other racial minority candidates preference over equally qualified white candidates in getting a job? [Strongly support/Support/Neither support nor oppose/Oppose/Strongly oppose]