

Race and Redistribution in the United States: An Experimental Analysis*

Jesper Akesson, Robert W. Hahn,
Robert D. Metcalfe and Itzhak Rasooly[†]

February 12, 2022

Abstract

It is often argued that many White Americans oppose redistributive welfare policies because of animosity toward Black Americans. While a host of observational studies lend credence to this view, it has not yet been rigorously tested using the tools of randomized inference. In this study, we do exactly this by conducting two incentive-compatible experiments ($n = 9,775$) in which different participants are randomly given different information about the share of welfare recipients who identify as Black and White. Our empirical analysis yields four main findings. First, most respondents greatly overestimate the share of welfare recipients who are Black. Second, these exaggerated beliefs reduce support for welfare among White respondents, a causal claim that we establish using treatment assignment as an instrument for beliefs. Third, just making White participants think about the racial composition of welfare recipients—holding constant their actual beliefs—reduces their support for welfare. Fourth, providing accurate information about the racial composition of welfare recipients does not influence support for welfare (relative to not receiving any information).

Keywords: field experiment, beliefs

JEL codes: D90, C93

*We would like to thank Filippo Muzi-Falconi for excellent research assistance. University of Oxford IRB: SOGE1A2020-185; AEA registry number AEARCTR-0008453.

[†]Akesson, The Behaviouralist; Hahn, University of Oxford, Carnegie Mellon University, and Technology Policy Institute; Metcalfe, University of Southern California and NBER; Rasooly, University of Oxford.

1 Introduction

Scholars have suggested that the racialized preferences in combination with beliefs about the ethnic composition of welfare recipients may limit support for U.S. welfare and redistribution (see, e.g., [Quadagno et al. \(1994\)](#); [Lee and Roemer \(2006\)](#); [Gilens \(2009\)](#)). As [Alesina et al. \(2001\)](#) put it, “Racial animosity in the US makes redistribution to the poor, who are disproportionately Black, unappealing to many voters.” This hypothesis has been supported by substantial observational evidence. For example, [Luttmer \(2001\)](#) finds an association between individual support for welfare and the fraction of local welfare recipients who share their ethnicity; similarly, [Alesina et al. \(2001\)](#) note that the US spends much less on welfare than other more ethnically homogeneous countries. However, to our knowledge, this hypothesis has not been directly tested using the tools of randomized inference.

In this paper, we provide such a test using two large-scale survey experiments. In total, the experiments have 9,775 participants, 7,929 of whom are White and 1,846 of whom are Black. In both experiments, individuals are randomly allocated to a control group that receives no information, a ‘low’ signal group, or a ‘high’ signal group. The low signal group receives an estimate suggesting that a relatively low share of welfare recipients are Black (20% in experiment 1, and 8% in experiment 2). In contrast, the high signal group receives an estimate suggesting a relatively high share of welfare recipients are Black (26% in experiment 1, and 52% in experiment 2).¹ We then elicit beliefs about the ethnic composition of welfare recipients using an incentive-compatible procedure. We also measure the support for welfare by asking participants if they would like to donate money to nonprofit organizations that work to increase or decrease welfare spending in the U.S. By using treatment assignment as an instrument for individual beliefs about the ethnic composition of welfare recipients, it is possible to obtain estimates of the effect of such beliefs on the support for welfare using two-stage least-squares regression (2SLS).

The experiments also contain other features aimed at understanding the link between race and welfare support. For example, we embed a priming experiment within the control group of our second experiment by randomly varying whether we measure support for welfare before or after eliciting their beliefs about the racial composition of welfare recipients. In addition, we ask participants whether they take race into account when deciding whether to support welfare spending, as well as various questions about their attitudes toward welfare

¹We do not deceive participants. Instead, we present them with estimates based on different sampling time-frames of the Survey of Income and Program Participation (2018). More specifically, we obtain variation in the estimates by measuring the share that are on welfare for different years and months of birth. We explicitly tell participants that the estimates they view are based on a particular sub-sample.

recipients.

Our empirical analysis yields four main findings. Our first finding is that participants dramatically overestimate the fraction of welfare recipients who are Black.² More specifically, the average White respondent in our sample estimates that 38% of welfare recipients are Black, and the average Black respondent estimates the figure at 35%. These numbers greatly exceed the true value, which is approximately 21% (Survey of Income and Participation, 2018). The finding that people overestimate the share of welfare recipients that are Black matches up with prior surveys on this issue (e.g., HuffPost/YouGov (2018)).

Second, we examine how beliefs about the share of welfare recipients who are Black influence support for welfare. To do this, we exploit the exogenous variation in beliefs induced by our treatments. Specifically, we use the (randomized) treatment assignment as an instrument for beliefs about the ethnic composition of welfare recipients (while excluding the control group from our analysis). We can do this since the only channel through which treatment assignment can plausibly influence support for welfare is via beliefs about ethnic composition, meaning that our instrument is plausibly valid. This leads to our second finding: higher beliefs about the share of welfare recipients who are Black reduces White respondents' support for welfare. For example, in the first experiment, we find that a 1 percentage point increase in the share that are Black leads to a 0.7 percentage point increase in the share that donate to the anti-welfare nonprofit ($p = 0.023$). Indeed, a 10 percentage point shift in beliefs has as large an effect on the share that donate to the anti-welfare nonprofit as the difference in donation rates between conservative and liberal respondents. This result holds for a variety of White subgroups and is robust across different specifications.

In contrast the result for White respondents, we do not find a significant relationship between beliefs about the ethnic composition of welfare recipients and support for welfare among Black respondents. This may simply be because our study is underpowered to detect such effects amongst the smaller subsample of Black respondents. However, it is also consistent with the possibility that Black Americans, unlike White Americans, do not substantially take race in account when forming attitudes towards welfare.

We also discuss mechanisms that might explain our main result (i.e., that White Americans support welfare less when they believe that the share of welfare recipients that is Black is higher). In some sense, the explanations for this is transparent: White Americans (on average) prefer welfare spending that goes to White as opposed to Black Americans.

²In our study, we define welfare as the following four programs: Temporary Assistance for Needy Families (TANF), Medicaid, the Supplemental Nutrition Assistance Program (SNAP), Housing Assistance, and Supplemental Security Income (SSI).

Nonetheless, one can ask what in turn explains this fact. We are able to shed some light on this issue: we find evidence that higher beliefs about the share of welfare recipients who are Black lead to lower beliefs about the perceived worthiness of welfare recipients, which should in turn reduce support for welfare (Fong and Luttmer, 2009a). However, one should be careful not to over-interpret this finding, not least because individuals may use claims about ‘deservingness’ as an excuse for underlying racial animus Bursztyn et al. (2020b).³

Next, we turn our attention to the priming experiment, which yields our fourth finding: simply making White respondents think about the racial composition of welfare recipients makes them less supportive of welfare. More specifically, we find that answering the question about the share of respondents that are from different racial/ethnic groups before being asked whether they would donate—as opposed to after—reduces the share that donate to the pro-welfare nonprofit by 4.9 percentage points ($p = 0.045$). This effect is fairly large relative to the effect of beliefs. For comparison, increasing participants’ perception about the share of welfare recipients that are Black by 1 percentage point decreases the share that donate to the pro-welfare nonprofit by 0.1 percentage points in experiment 2 ($p = 0.037$).

Finally, our fourth finding is that providing individuals with accurate information about the racial composition does not seem to alter their attitudes towards welfare. This might seem surprising since providing such information shifts beliefs about the prevalence of Black welfare recipients downwards, which in turn might be expected to increase support for welfare (in light of our earlier result). However, one should recognize that information provision also makes the issue of race more salient, which in turn could be expected to decrease support for welfare. Given these conflicting effects, it is not unexpected that information provision does not in itself have a large effect on welfare attitudes in either direction.

Our results build on a number of important literatures. Most importantly, there is a large literature examining whether beliefs about the ethnic distribution of welfare recipients alter support for welfare in a US context. Gilens (1995), Gilens (1996), Ribar and Wilhelm (1999), Alesina et al. (2001), Luttmer (2001) and Alesina and Glaeser (2004) all use observational data to make the case that believing that a greater proportion of welfare recipients are Black leads to less support for welfare spending. On a more theoretical level, Lee and Roemer (2006), Lee et al. (2006), Alesina and Stantcheva (2020) show how such beliefs ought to determine both welfare support and therefore equilibrium outcomes. Our study builds on

³For this result, we estimate our 2SLS specification but have worthiness of welfare recipients and efficacy of the welfare state as the dependent variables (instead of giving to the charities). We estimate that higher signals leads to lower worthiness of welfare recipients. We did not estimate these mechanisms by changing the actual treatments to identify the discrimination through worthiness or racial animus, as in List (2004); Flory et al. (2015).

this literature by providing the first experimental (and therefore causally well-identified) evidence on this important issue.⁴

More broadly, our findings are also consistent with the literature on the impact of immigration on preferences for redistribution. A number of studies find that an influx of immigrants can dampen support for redistribution in a variety of contexts (Dahlberg et al., 2012; Kraus et al., 2019). In addition, Alesina et al. (2018) find that just priming individuals to think about immigration reduces their support for redistribution. Both of these findings are echoed by our results.

Moreover, our study contributes to a growing literature that examines whether information provision can alter behavior and attitudes (see Haaland et al. (2020) for a review). Especially related are the studies that provide race-related information, such as estimates of the Black/White wealth gap (Onyeador et al., 2021; Alesina et al., 2021; Callaghan et al., 2021)⁵, the amount of discrimination faced by Black Americans in the labor market (Haaland and Roth, 2021), and information about White privilege (Phillips and Lowery, 2015). Indeed, our fourth finding—that providing accurate information fails to change behavior—echoes the null result obtained by Alesina et al. (2018).

Finally, at the highest level, our study contributes to the wider literature on racial disparities within the US. Studies on this topic have documented a wide array of inequalities in social, legal and economic outcomes between White and Black Americans (Chiteji and Stafford, 1999; Barsky et al., 2002; Gittleman and Wolff, 2004; Altonji and Doraszelski, 2005; Charles and Hurst, 2002; Rothstein, 2017; Derenoncourt et al., 2021). Insofar as our findings map into welfare policy (discussed later), they may help go some way to explaining a portion of these disparities.

The remainder of our paper is structured as follows. Section 2 outlines the design of our two

⁴There are also a number of papers that although highly related, do not directly study the impact of beliefs about the ethnic distribution of welfare recipients on welfare support. For example, O'Brien (2017) conducts a survey experiment and finds that White American support for taxation is influenced by the share of the population that is Hispanic. Wetts and Willer (2018) conduct two survey experiments and show that White Americans are less likely to support welfare if the programs are framed as primarily benefiting Black Americans, and that they are also less supportive of welfare when presented with information suggesting that Whites' status as a majority group is rapidly coming under threat. Bobo and Kluegel (1993) also find observational evidence suggesting that White Americans are opposed to race-targeted welfare policies. Finally, Gilens (1996) and Harell et al. (2016) conduct hypothetical vignette experiments, which suggest that providing cues about the race of particular individuals influence the amount of support that subjects deem they should receive (see also Fong and Luttmer (2009a), Fong and Luttmer (2009b) and Gross and Wronski (2021)).

⁵The fact that individuals have misconceived beliefs about race-related issues is also picked up by studies on perceptions of racial equality in the US—see Brodish et al. (2008); Eibach and Keegan (2006); Kraus et al. (2017, 2019).

experiments, and Section 3 provides our central results. Finally, Section 4 concludes with a discussion of the policy implications of our results and directions for future research.

2 Experimental design

We conducted two experiments, the first ($n = 5,793$) in January 2021 and the second ($n = 3,982$) in October 2021. Participants were recruited via Prolific Academic and were paid a flat fee of \$1.60 in exchange for their participation (they were also entered into various lotteries, as outlined below).⁶ The experiments took place within two Qualtrics surveys, both of which took an average of 13 minutes to complete. We also recorded the time taken by respondents on all important questions. While the experiments share many similarities, there are some key differences that we describe below.

2.1 Experiment 1

The first experiment began by asking respondents some standard demographic questions, including their age, state of residence, educational attainment, gender, race/ethnicity, household income, political affiliation, and whether they had ever been on welfare.⁷ In the course of asking respondents whether they had been ‘on welfare’, we explicitly defined what we meant by the term ‘welfare’ for the purposes of the survey. More specifically, we told respondents that they would take welfare to refer to any of the following programs:⁸

1. Temporary Assistance for Needy Families (TANF)
2. Medicaid
3. Supplemental Nutrition Assistance Program
4. Housing Assistance
5. Supplementary Security Income

⁶More information about Prolific Academic can be found at <https://www.prolific.co/>. Peer et al. (2017) show that participants recruited via Prolific Academic are less dishonest, are less likely to fail attention checks, and produce higher quality data than participants recruited via other comparable online research platforms.

⁷In addition, we also asked participants about their news consumption, how much they believe that the US government spends on welfare, and whether they voted in the 2020 Presidential Election.

⁸As an empirical matter, it is unclear whether the extent to which welfare attitudes depend on race depend on the precise programs included in the definition. Nonetheless, we thought that an explicit definition might be helpful, not least because it allows us to calculate whether individual perceptions about welfare are accurate. Our definition of welfare includes the five largest means-tested redistributive programs (excluding EITC, as this program has a work requirement).

Before moving to the main treatment, we exposed respondents to a ‘strategic attention check’ (following [Alesina et al. \(2018\)](#)). To do this, we asked respondents whether we should use their responses, or instead whether they should discard their responses since they had not devoted their full attention to the questions so far. The main aim of this question was to prompt respondents to pay attention in the next (and more important) section of the survey.

We then elicited respondents’ beliefs about the share of welfare recipients who identify as White, Black, or as belonging to other races and ethnicities. We also asked what proportion of welfare recipients they thought identify as Spanish/Hispanic/Latino. We incentivized accurate answers to these questions by telling respondents that answers within 2 percentage points of the truth would be rewarded with entry into a lottery for \$100. We closed the belief elicitation section by asking respondents about their degree of confidence in their estimates.

Having elicited respondents’ beliefs, we then randomly assigned respondents into either the ‘high’ treatment group, the ‘low treatment’ group, or a control group. The high treatment group were told that:

Estimates from 2017 suggest that, out of every 100 American adults on welfare,*

- *63 were White*
- *26 were Black*
- *11 belonged to other ethnic groups*

**Please note that these estimates were obtained using the Survey of Income and Program Participation (2018). The statistics were computed for individuals who were born in the month of October (in any year) and may thus not be fully representative of the overall population. See <https://www.census.gov/programs-surveys/sipp.html> for more information.*

Those in the low treatment group were also given estimates about the racial/ethnic distribution of welfare recipients. However, our estimates were obtained from those born in September (in any year), and yielded the following distribution: 68% White, 20% Black, and 12% belonging to other groups. Finally, the control group received no such information. We stress that, although we provided different information to members of the different treatment groups, all information provided was fully accurate and the nature of the estimates was disclosed.

In an effort to ensure that participants processed the information provided, we next asked them to recall the estimates that they had just been shown. Respondents were then told if their answer had been correct, and were then shown the treatment estimates for a second

time. While the main motivation of this quiz was to encourage respondents to further internalize the estimates, the answers to the quiz may also have provide further information about the attentiveness of particular subjects (as discussed below).⁹

Immediately after the treatments, we once again elicited respondents’ beliefs about the ethnic distribution of welfare recipients. To do this, we told respondents that, while the estimates which they had been presented were from 2017, they should instead “think about now”. They were then asked, out of every hundred American adults on welfare, how many are Black, White, and neither. As before, we incentivized correct answers (within 2 percentage points of the truth) and asked respondents how confident they were in their answers (on a 5-point Likert scale).

After re-eliciting beliefs, we recorded our main outcome of interest, namely whether respondents supported welfare. To measure this in an incentive-compatible way, we first told respondents that they would automatically be enrolled into a lottery for \$100. We then asked them whether they would like to donate their potential winnings to either of two non-profit organizations, one chosen to be ‘pro-welfare’ and the other chosen to be ‘anti-welfare’. Respondents were given the following information about the non-profits:

1. The Center on Budget and Policy Priorities is a progressive American organization (think tank) that works to ensure that policymakers consider the needs of low-income people. Many of these people receive food stamps, housing assistance or other forms of welfare assistance.
2. The Foundation for Government Accountability is an organization (think tank) that focuses on welfare and health care reform. Many of the policies proposed by this think tank would have the impact of reducing federal welfare spending.

We used an incentive-compatible outcome (donations) in order to minimize any possible experimenter demand effects and ensure that subjects were attentive when providing answers. This approach has also been taken in other experimental studies on political attitudes. For example, [Bursztyn et al. \(2020a\)](#) use donations to a xenophobic organization as their main outcome of interest when studying how social norms influence racist expression. Somewhat closer to our setting, [Alesina et al. \(2018\)](#) ask respondents if they would like to donate prospective lottery winnings to charities supporting low-income adults or children in an attempt to capture attitudes toward redistribution.

⁹Note that participants in the control group were given this quiz as they did not receive any treatment information. For this same reason, we also did not re-elicite their beliefs about the racial/ethnic distribution of welfare recipients.

Having measured respondents’ attitudes toward welfare, we then asked them two questions aimed at detecting the mechanism through which particular effects might be operating. Past studies (see, e.g., [Cook and Barrett \(1992\)](#), [Fong \(2001\)](#), [Henry et al. \(2004\)](#)) have identified perceptions about both the ‘efficacy’ of welfare spending (at alleviating poverty) and the ‘deservedness’ of welfare recipients as key determinants of support for welfare. They have also identified differential beliefs about the deservingness of Black as opposed to White welfare recipients as well as about how welfare relatively benefits these groups (see, e.g., [Fong and Luttmer \(2009a\)](#), [Gilens \(1996\)](#)). In light of these findings, we asked respondents whether they thought that “welfare programs help lift Americans out of poverty” to capture an “efficacy” channel; and whether they thought that “people who receive welfare are poor through no fault of their own” to capture a “deservingness” channel.¹⁰

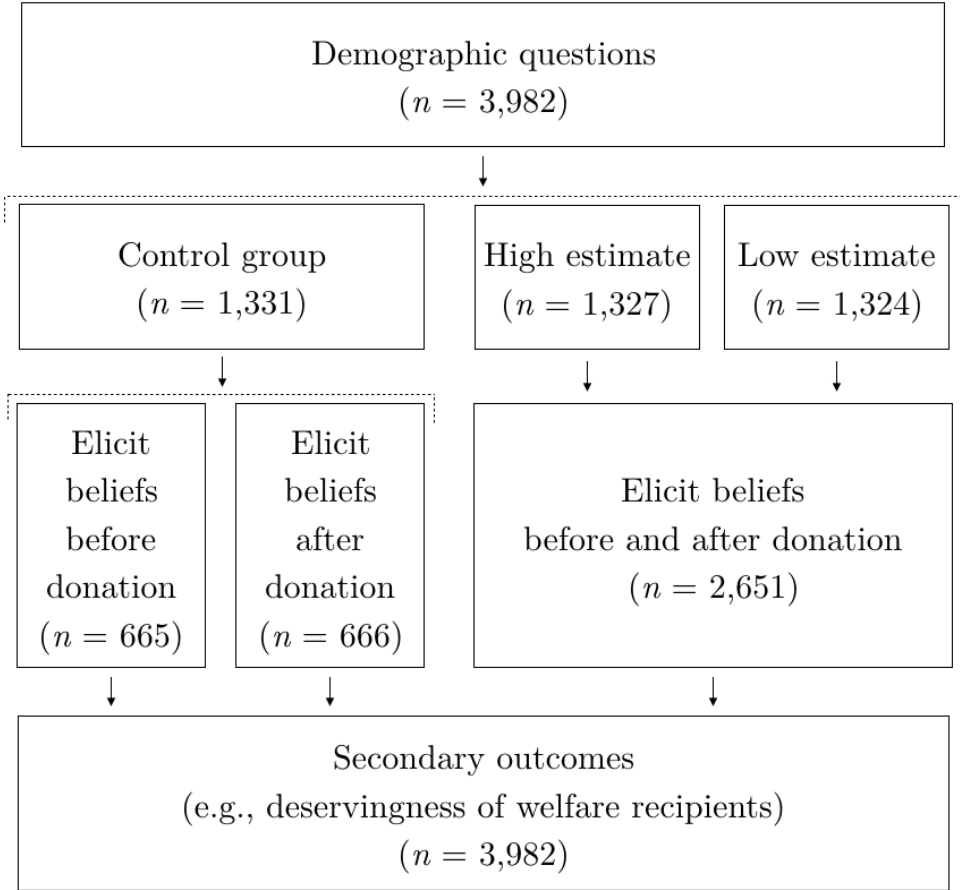
We concluded the experiment by asking respondents about their likelihood of winning the lottery to ensure the validity of our incentivization. We also measured their ‘implied beliefs’ about the share of welfare recipients that are from different racial/ethnic groups. To do this, we asked what share of different groups they believed were on welfare and what share of the population belong to different groups (allowing for a calculation of ‘implied beliefs’ by Bayes rule). Finally, we asked an open-ended question on what they thought the study was about. This question allowed us to drop those who understood that the survey was about racial attitudes (as a robustness check).

2.2 Experiment 2

The structure of the second experiment broadly mirrored that of the first (see [Figure 1](#) for an outline of the experimental structure). However, there were several important differences. First, we dramatically increased the difference in the estimates with which the two treatment groups were provided in order to ensure that our results are not sensitive to the exact parameter values used in the first experiment. Rather than providing information on welfare recipients born in different months (Experiment 1), we provided information on welfare recipients born in different months of a particular year in Experiment 2 (December 1986 for the ‘low’ group and April 1987 for the ‘high’ group). This generated an 8% Black, 75% White, and 17% ‘other’ distribution in the ‘low’ treatment group; and a 52% Black, 31% White, and 17% ‘other’ distribution in the ‘high’ treatment group. As in the first experiment, we clearly informed participants how the estimates had been generated, and that they might be unrepresentative of the population as a whole.

¹⁰See the survey for the full list of the mechanism questions.

Figure 1: Experimental design (experiment 2)



Notes. In this figure, we present the experimental procedure for experiment 2. The structure of experiment 1 is very similar, with the main difference being that we do not randomize those in the control condition into two groups.

Second, we recorded the extent to which participants displayed ‘social desirability bias’ (i.e., a tendency to say and do things in order to confirm with the relevant group) prior to presenting the treatment information. To do this, we asked participants questions from the Marlowe-Crowne Social Desirability Scale (Crowne and Marlowe, 1960). We elicited this information to allow us to conduct a robustness check where we drop those who exhibit high degrees of social desirability bias as these individuals may be less likely to reveal their true beliefs and attitudes—particularly if they realized what the survey was about. This approach has been used in a number of other experimental studies, such as Dhar et al. (2018)).

Third, similar to Alesina et al. (2018), we embedded a priming experiment in our survey. More specifically, those who were randomly assigned to the control group were further randomized into one of two groups: 1) those asked whether they would like to donate money to

the non-profits before being asked about the share of welfare recipients that belong to different ethnic/racial groups, and 2) those asked if they would like to donate money after being asked about the share of welfare recipients that belong to different ethnic/racial groups. In other words, we randomized whether the question about ethnicity came before or after the questions about welfare support. This allowed us to investigate whether being primed to think about race might influence support for welfare.

Finally, we included two additional questions at the end of the survey that provide us with descriptive evidence about the relationship between people’s beliefs about the share of welfare recipients that are Black and their support for welfare. The first question asked whether respondents took race into account when deciding whether to donate money to one of the two non-profits. The second question asked whether participants have thought about the share of welfare recipients that belong to different ethnic/racial groups prior to taking the survey. We asked these questions to further understand whether these beliefs about race matter for expressed welfare support both within the survey and beyond it.

2.3 Data and sample

In total, we sampled 9,775 participants, 5,793 of whom took part in the first experiment and 3,982 of whom took part in the second experiment. All participants were white, with the exception of 1,846 black participants who took part in the first experiment.¹¹ While we made some effort to sample representatively, the sample was not altogether balanced on common demographic characteristics. Notably, the sample was disproportionately female (59% female and 66% female in experiments 1 and 2 respectively) and disproportionately liberal (58%) in experiment 1—see Tables A1 and A2 for descriptive statistics. As a result, we checked that our results were robust to re-weighting.

3 Results


3.1 Beliefs about race and welfare

We begin by examining the accuracy of respondent beliefs about the racial/ethnic distribution of welfare recipients. Figure 2 plots respondent beliefs about the share of welfare recipients who are Black and White, pooled across the two experiments. Respondents estimated on average that 37% of welfare recipients are Black, 39% are White, that 25% of

¹¹We primarily recruited White participants since our main hypothesis (i.e., that White individuals dislike redistributing to Black participants) concerns this group.

welfare recipients are neither, and that 27% of welfare recipients are Hispanic. Participants are somewhat confident in their beliefs, with only 37% of respondents stating that they are ‘very unconfident’ or ‘somewhat confident’ that their estimate is correct. These average estimates, however, are quite far from the truth. For example, according to the Survey of Income and Program Participation, 21% of welfare recipients are Black. Therefore, the average respondent substantially overestimates the share of welfare recipients who are Black (with 86% overestimating this figure). This finding is consistent with prior survey evidence on this issue (Gallup, ???).

Figure 2: Beliefs about the ethnic distribution of welfare recipients



Notes. In this figure, we present the distribution of participants’ pre-treatment beliefs (pooled across both experiments) about the share of welfare recipients who are White and Black. The Y-axis shows the share of participants who held a particular belief, and the X-axis shows the particular point estimate.

Next, we examine which groups are most likely to overestimate the share of welfare recipients who are Black. To do this, we calculate the fraction of respondents in various subgroups who estimate the share at least 26 percentage points (i.e., an overestimate of at least 5 percentage points). Table A3 displays the results (see also Table A4 for a linear probability model that

reports similar findings). As can be seen, conservatives and moderates are significantly more likely to overestimate the share of welfare recipients who are Black than liberals. Similarly, women are more likely to overestimate the figure than men, and White respondents are more likely to overestimate the share than Black respondents. Strikingly, however, each of these subgroups still overestimates the share substantially, even including Black respondents who on average estimate the figure at 35%. This underscores the robustness of this phenomenon.

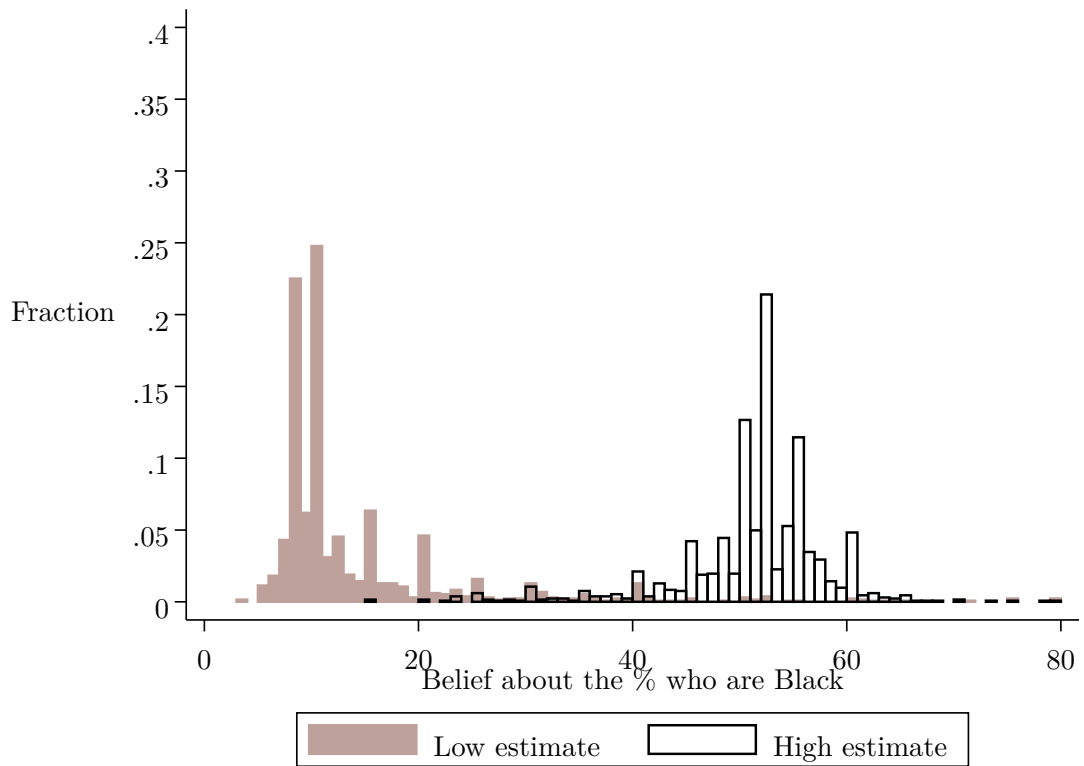
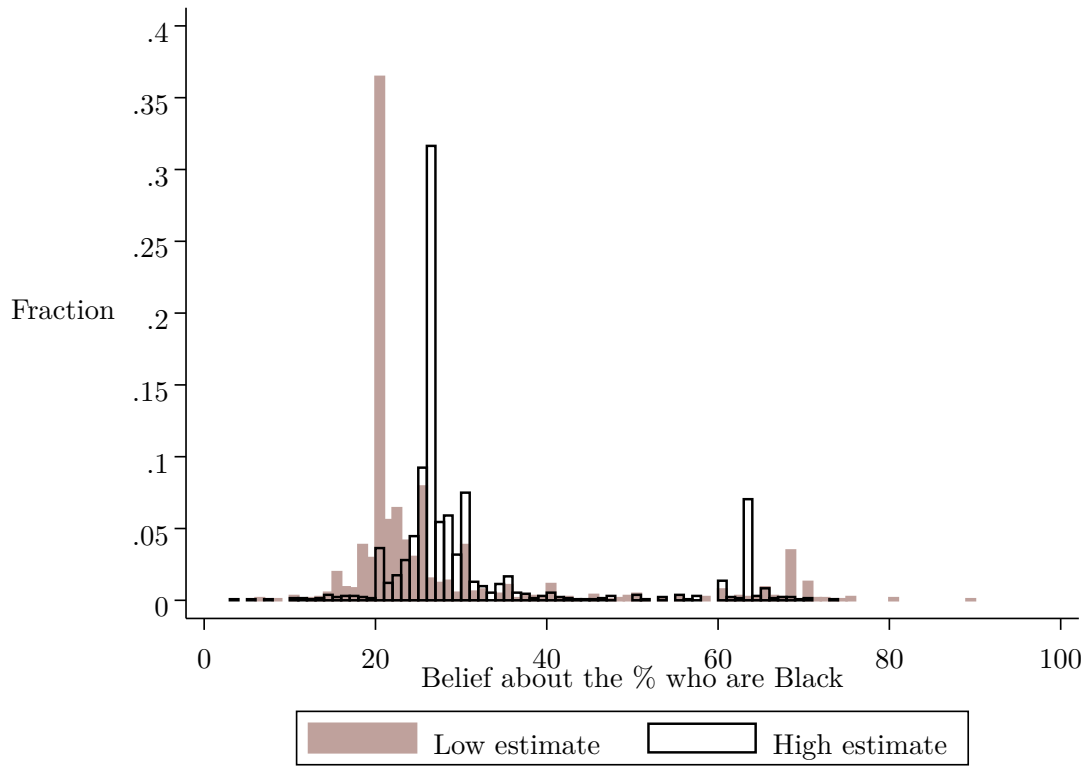
In principle, there are a number of reasons why we might observe these systematic overestimates. While we are unable to examine all factors that contribute to these beliefs, three potentially important factors—suggested by Bayes rule—are beliefs about the share of Black people who are welfare recipients, beliefs about the share of the US population that is Black, and the share of the US population that is on welfare. While we find that respondents do not exactly use Bayes rule, a regression of our outcome on these three variables suggests that they do influence our outcome in the expected direction: for example, increasing beliefs about the share of the population that is Black by one percentage point is associated with a 0.242 percentage point increase in the perceived share of welfare recipients who are Black (see Table A5). Moreover, we find that respondents greatly overestimate this statistic, putting the share of the US population that is Black at 27%. This seems likely to be a contributor to the widespread overestimation.¹²

3.2 Updating beliefs

We now examine how our treatments influence beliefs about the share of welfare recipients who are Black, only focusing on White respondents as our main hypothesis (i.e., that White respondents dislike redistribution to Black respondents) concerns this subgroup. Figure 3 displays the distribution of beliefs in the high and low treatment groups of both experiments (the upper panel for experiment 1 and the lower panel for experiment 2). As can be seen, both experiments succeeded in shifting the distribution of beliefs in the expected way. In experiment 1, the distributions are centered around 20% and 26% (i.e., the exact estimates with which subjects were presented in their respective treatments). In experiment 2, the distributions are centered around 8% and 52%, which are again the estimates from the two treatments. While the differential assignment to the treatments shifted beliefs in the expected way, there is some clustering around the provided estimates. In part, this may be because subjects were presented with information from a sub-sample from 2017, and therefore chose to make some revisions when estimating the figure for the population in 2020.

¹²See [Augenblick and Rabin \(2021\)](#) and associated references for a discussion of insufficient sensitivity to the various components of Bayes rule.

Figure 3: Belief updating



Notes. The top panel presents the distribution of beliefs about the share of welfare recipients who are Black in the high- and low-information groups, respectively, in the first experiment. The bottom panel does the same for the second experiment.

As discussed earlier, we manipulate beliefs in this way so that we can use treatment assignment as an instrument when estimating the effect of beliefs on attitudes towards welfare. To check that our instruments are informative, we now regress beliefs on treatment assignment in both experiments (omitting the control groups from the analysis). As can be seen from Table A6, allocating individuals to the high as opposed to low treatment group increases beliefs by 4.16 percentage points on average in experiment 1 ($p < 0.0001$). In experiment 2, allocating individuals to the high as opposed to low treatment group increases their beliefs on average by 36.9 percentage points. Since the F statistics are large in both cases (65 and 11253 in experiments 1 and 2 respectively), we can confidently conclude that our instruments are informative.

Next, we examine whether our instruments are valid. In other words, we investigate whether our instrument (treatment assignment) can influence our outcome (donations) through a variable other than beliefs (by which we specifically mean ‘point estimates’). We can only think of two ways in which our instruments might have influenced relevant factors other than beliefs: 1) they might influence the salience of beliefs; and 2) they might influence the confidence with which beliefs are held. Since both treatment groups are provided with similarly formatted information, we can reject the first possibility: both groups should be equally primed to consider the information as salient. To investigate the second possibility, we estimate the effect of treatment assignment on stated confidence. As Table A7 shows, we do not find a significant difference in the confidence with which beliefs are held between those allocated to the high and low groups in both experiments.¹³ We can therefore conclude that any confidence effect is unlikely to be an issue for biasing our key coefficient estimates.¹⁴

3.3 The effect of beliefs on support for welfare

Having argued that our instruments are informative and valid, we now turn to the instrumental variables analysis. In other words, we examine if beliefs about the share of welfare recipients who are Black influence White respondents’ support for welfare, using treatment assignment as an instrument for beliefs (and excluding the control group). Table 1 displays

¹³As a robustness check, we control for confidence when conducting the instrumental variable analyses and we do not find that this changes our estimates appreciably. We also examine the effects of the treatment using different ways of coding the confidence outcome (e.g., treating it as a binary or continuous variable). Please see Section 3.3 for more information.

¹⁴Note that these points would *not* apply to a comparison of either of our treatment groups with the control group. First, members of the treatment groups turn out to be substantially more confident in their estimates than members of the control group, which is exactly what one would expect given that they (unlike the control group) have been given information. Second, and for this reason, members of the treatment groups may have been more primed to think about race than members of the control group. For this reason, we omit the control group from our main analysis.

both our results along with the associated Intent to Treat (ITT) estimates.

Table 1: Main results

	Experiment 1			
	Pro-welfare	Pro-welfare	Anti-welfare	Anti-welfare
High estimate	-0.0262 (0.0185)		0.0298** (0.0127)	
Belief % Black		-0.00628 (0.00449)		0.00716** (0.00315)
Constant	0.358*** (0.0132)	0.526*** (0.130)	0.106*** (0.00847)	-0.0852 (0.0904)
<i>n</i>	2,646	2,646	2,646	2,646
	Experiment 2			
	Pro-welfare	Pro-welfare	Anti-welfare	Anti-welfare
High estimate	-0.0361** (0.0173)		0.000559 (0.0109)	
Belief % Black		-0.000979** (0.000469)		1.52e-05 (0.000296)
Constant	0.292*** (0.0125)	0.305*** (0.0178)	0.0861*** (0.00771)	0.0859*** (0.0110)
<i>n</i>	2,651	2,651	2,651	2,651

Notes. In this table, we LATE and ITT estimates. We exclude the control group from the analysis. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$.

In experiment 1, we find that beliefs have a significant effect on the share who donate to the anti-welfare organization. More specifically, for every one percentage point increase in beliefs about the share who are Black, participants become 0.7 percentage points more likely to donate to the anti-welfare organization ($p = 0.02$). We also find evidence suggesting that higher beliefs about the share who are Black make participants less likely to donate to the pro-welfare organization, although the estimated coefficient is not statistically different from zero.

In experiment 2, the findings are similar. We find a significant effect of beliefs on the share who donate to the pro-welfare organization: for every percentage point increase about the share who are Black, participants become 0.1 percentage points less likely to donate to the organization. Moreover, while we do not find a significant relationship between beliefs and the share who donate to the anti-welfare organization, we do find that the estimated coefficient is positive. In summary, then, all four coefficient estimates (from both experiments) suggest that higher beliefs about the share who are Black reduce support for welfare, and two of these four estimates are significant.

We can also examine this issue after pooling the data from both experiments, as this affords

us greater statistical power. To do this, we now instrument for beliefs using a dummy variable indicating treatment assignment (high or low), a dummy variable indicating which experiment participants took part in (experiment 1 or 2), and the interaction of these two variables. As before, all instruments are clearly informative. Moreover, in light of the previous discussion, the treatment assignment instrument is plausibly valid.¹⁵ However, since participants were not assigned to experiments randomly, there is a possibility that this instrument influences our outcome of interest because different types of people took part in the two experiments, rather than just affecting it through beliefs. We therefore control for all observables when conducting our pooled analyses.

As recorded in Table A8, the pooled results closely resemble those of experiment 2. As before, we observe that increasing beliefs about the share of welfare recipients who are Black significantly decreases the share who donate to the pro-welfare organization ($p < 0.01$). The estimate is that a one percentage point increase in beliefs decreases the share who donate by 0.12 percentage points, an effect that is quantitatively very similar to our previous estimate. As in experiment 2, we also do not obtain a statistically significant effect on donations to the anti-welfare organization.

The fact that we obtain qualitatively similar effects in two separate experiments suggests that the phenomenon that we uncover may generalize. To further study this issue, however, we now conduct several sets of sensitivity analyses. The first set of checks we conduct involves dropping apparently less attentive participants. We begin by dropping participants who completed the survey in less than t minutes, for various values of t , on the grounds that such participants may have failed to pay proper attention to the questions (see Figure X). We also drop participants who spent less than t seconds on the treatment screens, and who spent less than t seconds when providing us with their belief estimates (again doing this for various values of t). Finally, we drop participants who failed to answer the post-treatment quiz correctly (i.e., the question where we asked participants to recite the treatment information), and the very small number of participants who stated that the share of welfare recipients who are Black is 0% or 100%. In general, our estimated coefficients remain quantitatively similar and statistically significant, despite the reduced sample size.

The second set of checks we conduct involves excluding participants who may have failed to provide fully truthful and reflective answers due to experimenter demand or a lack of incentivization. We do this by dropping participants who thought that it was unlikely that

¹⁵The treatments could, in theory, influence the confidence with which people hold their beliefs. If this is the case, our instruments would no longer be valid. However, we also measure people’s confidence, and we do not find that the treatment has an effect on this variable.

they would win the lotteries, on various definitions of ‘unlikely’. We also exclude participants who displayed a high degree of ‘social desirability bias’ (as measured by the Marlow-Crowne scale), and those who understood what the study was about, as judged by their answers to the open-ended debriefing question at the end of the survey. Again, our estimated coefficients remain quantitatively similar, but also become more statistically significant.

We also conduct some additional robustness checks. First, we re-weight our data so that it matches the joint distribution of gender, age, and income for the US population. Second, we conduct the analyses while controlling (and instrumenting) for the confidence with which beliefs are held. Third, we re-estimate the regressions using participants’ implied beliefs about the share of welfare recipients who are Black (we obtain this using their estimates of the share of the US population that is Black, the share of Black people who are on welfare, and the share of the US population that is on welfare). Finally, we re-estimate our regressions including the control group in the analysis. As in the previous cases, none of these exercises alter our substantive conclusions (see Table [A9](#) and [A10](#) for the results of all robustness checks).

We can also examine whether the effects we uncover can be detected in the descriptive data. To investigate this, we regress donations on participants’ beliefs about the share of welfare recipients who are Black, attempting to control for all relevant confounds in the hope that this might lend our regression a causal interpretation. We also utilize different measures of beliefs (i.e., those elicited prior to treatment assignment and those elicited afterwards) when conducting these regressions. Strikingly, we find that the simple associations between beliefs and support for welfare are close to identical to the estimates that we obtain in our instrumental variable analyses (see Table [A11](#)).

Aside from testing the robustness of our results, it is also important to understand how they should be interpreted. The most natural interpretation is that our treatments have an effect simply by increasing the share of welfare recipients deemed to be Black, and decreasing the share deemed to be White. However, other alternatives are possible: it might be that increasing beliefs about the share who are Black also alters beliefs about the share on welfare who are neither White nor Black (i.e., those who belong to ‘other’ ethnic group). Similarly, our treatment may also shift beliefs about the share of welfare recipients who are Hispanic/Latino. To examine this issue, we measure the effects of treatment assignment on participants’ beliefs about the share of welfare recipients who are White, belong to ‘other’ ethnic/racial groups, or who are Hispanic/Latino. As can be seen in Table [2](#), the treatments do not influence beliefs about the share of welfare recipients who belong to ‘other ethnic/racial groups’, and similarly do not meaningfully influence the share of welfare re-

Table 2: Treatment effect on other beliefs

Variables	Belief % White (exp. 1)	Belief % Other Other (exp. 1)	Belief % Hispanic (exp. 1)	Belief % White (exp. 2)	Belief % Other (exp. 2)	Belief % Hispanic (exp. 2)
High estimate	-3.517*** (0.559)	-0.651*** (0.246)	-0.231 (0.659)	-37.07*** (0.395)	0.224 (0.243)	3.591*** (0.476)
Constant	59.94*** (0.419)	13.30*** (0.165)	20.31*** (0.479)	68.47*** (0.357)	17.56*** (0.180)	19.32*** (0.320)
n	2,646	2,646	2,646	2,651	2,651	2,651
R^2	0.015	0.003	0.000	0.769	0.000	0.021

Notes. In this table, we present treatment effects on participants’ beliefs about the share of welfare recipients that is White, belongs to Other ethnic/racial groups, or Hispanic/Latino. We exclude the control group in the analyses. Standard errors in parentheses. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$.

ipients who are Hispanic/Latino. This suggests that the most obvious interpretation is the right one: the treatments alter beliefs about the share of welfare recipients who are Black whilst simultaneously altering beliefs about the share who are (non-Hispanic) White.

3.4 Mechanisms

Here we discuss the mechanisms underlying our main result. In some sense, the mechanism is very transparent: if White Americans support welfare less when they believe that a higher share of welfare recipients are Black, this must be because they place less value on welfare spending that goes to Black Americans. Nonetheless, one can ask the further question of what explains this underlying preference. In practice, it is very difficult to adjudicate between different explanations (e.g., taste-based vs statistical theories), primarily because individuals might pretend to object to giving welfare on ‘statistical’ grounds (e.g., because welfare is less effective for some groups than others) even when their objections are rooted in racial animus (Burszty *et al.*, 2020b). However, we now attempt to shed some light on this issue.

We do this by examining the extent to which White people’s beliefs about the share of welfare recipients who are Black also influence the extent to which they agree with the following statements: “people who receive welfare are poor through no fault of their own” and “welfare programs help lift Americans out of poverty”. Table 3 presents the results. As we can see, we do not find a significant relationship between beliefs and agreement with these statements in either experiment, although our estimated coefficients are negative (i.e., participants are less likely to, for example, think that welfare is effective if the share of welfare recipients who are Black increases). We do, however, find a significant relationship

Table 3: Secondary outcomes

Variables	Deservingness (exp 1)	Deservingness (exp 2)	Deservingness (pooled)
Belief % Black	-0.00161 (0.00465)	-0.000824 (0.000520)	-0.000943** (0.000473)
Constant	0.574*** (0.135)	0.450*** (0.0195)	0.554*** (0.0298)
Controls	No	No	Yes
<i>n</i>	2,646	2,651	5,297
	Efficacy (exp 1)	Efficacy (exp 2)	Efficacy (pooled)
Belief % Black	0.00136 (0.00452)	-1.39e-05 (0.000523)	-0.000132 (0.000485)
Constant	0.588*** (0.131)	0.560*** (0.0195)	0.792*** (0.0300)
Controls	No	No	Yes
<i>n</i>	2,646	2,651	5,297

Notes. In this table, we present estimates of the effect of beliefs about the share of welfare recipients who are Black on beliefs about the deservingness of welfare recipients and the efficacy of welfare. We control for demographic characteristics when conducting the pooled analysis. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$.

between the share of welfare recipients who are Black and whether participants think that welfare recipients are poor through no fault of their own when pooling the data from the two experiments. This could be taken to suggest that a ‘deservingness’ channel underlies our main result, particularly as other studies have found a casual effect of perceived deservingness on charitable donations (c.f. [Fong and Luttmer \(2011\)](#)). That said, one should treat this with caution since statements about ‘deservingness’ may function as an insincere excuse for racist attitudes [Bursztyn et al. \(2020b\)](#).

3.5 Priming

The previous section studies how attitudes towards welfare depend on the precise beliefs individuals hold about the ethnic distribution of welfare recipients. Inspired by [Alesina et al. \(2021\)](#), we now study whether simply prompting White respondents to think about the share of welfare recipients who are from different races also changes their attitudes towards redistribution. To do so, we turn our attention to the second experiment, where we randomly varied whether the question about the ethnic distribution of welfare recipients came before or after the questions about welfare support (for those in the control group).

Table 4 shows that those who were ‘primed’ to think about race (i.e., those who were asked about the race/ethnicities of welfare recipients first) are 4.9 percentage points less likely to donate to the pro-welfare organization ($p = 0.047$). Echoing the findings of [Alesina et al. \(2021\)](#), this suggests that simply getting individuals to think about race can dampen their support for redistributive programs. This may be because they believe that welfare recipients are disproportionately likely to be Black (see Section 3.1) and dislike redistributing funds to Black Americans (see Section 3.3).¹⁶

3.6 The effects of correcting beliefs

We now examine the effect of providing accurate information about the ethnic distribution of welfare recipients. In light of our previous findings, it is unclear what effect this might be expected to have. On the one hand, since people generally overestimate the share of welfare recipients who are Black (Section 3.1), information provision should be expected to reduce beliefs about the share. This in turn might be expected to boost support for welfare (Section 3.3). On the other hand, providing such information also makes the issue of race more salient, which might be expected to decrease support for welfare (Section 3.5). It is therefore unclear what effect such information should have on net.

¹⁶As one would expect, we do not find that priming influences beliefs about the share of welfare recipients who are Black—suggesting that this is not the channel through which priming effects operate.

Table 4: Priming

Variables	Pro-welfare donation	Anti-welfare donation
Priming condition	0.0488** (0.0245)	-0.00325 (0.0150)
Constant	0.253*** (0.0174)	0.0828*** (0.0106)
n	1,330	1,330
R^2	0.003	0.000

Notes. In this table, we present eliciting White respondents' beliefs about the racial/ethnic composition of welfare recipients before or after asking them to donate to the pro- or anti-welfare organizations. We exclude the control group from the analysis. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$.

Table 5: The effects of correcting beliefs

Variables	Pro-welfare	Anti-welfare
High estimate	0.00596 (0.0187)	-0.0149 (0.0124)
Constant	0.353*** (0.0133)	0.121*** (0.00880)
n	2,627	2,627
R^2	0.000	0.001

Notes. In this table, we present the effect of presenting White respondents with accurate information about the share of welfare recipients that is Black. We exclude the control group from the analysis. * $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$.

To study this question, we compare support for welfare among those in the control group (i.e., those who received no information) to those in the ‘low’ treatment group in experiment 1 (as these individuals were presented with accurate information about the share of welfare recipients who are Black, namely that this figure is around 21%). As Table 5 reveals, we do not find any significant effect of the treatment on the share who donate to the pro- or anti-welfare organizations. This suggests that the ‘belief updating’ and ‘priming’ effects may be almost exactly canceling one another out.

3.7 Exploratory analyses

In this section, we conduct a number of exploratory analyses to shed more light on our key research questions and its underlying mechanisms. We first note that around 42% of the White respondents who donated to the anti-welfare organization said that they took race into

account when making this decision, which corroborates our main finding.¹⁷ Moreover, 48.5% of respondents (and 58.3% of those who donated against welfare) said that they have thought about the race/ethnicity of welfare recipients before, which suggests that the relevance of this factor extends beyond our survey.

Next, we regress support for welfare (measured by donations) on various factors that have been found to be important determinants in other studies. As Table A12 shows, all else being equal, we find that women are more supportive of welfare than men. We also find that liberals are also more pro-welfare than conservatives and moderates, that those who have been on welfare are more supportive than those who have not, and that the young are more supportive than the old. Reassuringly, all of these findings are consistent with prior work on this topic (see, e.g., Ashok et al. (2015)), a fact which further validates our measures of welfare support.

We can also compare the coefficient estimates obtained in our main analysis (see Section 3.3) with the associations just discussed. Strikingly, beliefs about the ethnic distribution of welfare recipients appear to be a rather important factor compared to these other determinants. For example, a 10 percentage point shift in beliefs has as large an effect on the share that donate to the anti-welfare nonprofit as the difference in donation rates between conservative and liberal respondents (this estimate was obtained from experiment 1). These comparisons suggest that the effects that we uncover are not just statistically but also quantitatively significant—and could therefore be an important determinant of welfare spending in the United States.

Finally, we examine if there is a relationship between beliefs about the ethnic/racial distribution of welfare recipients and support for welfare among Black respondents. To do this, we restrict our attention to experiment 1, which had close to 2,000 Black participants (see Table A13). While we find that the treatments update participants' beliefs in the expected way, we do not find that there is a significant relationship between beliefs and welfare support. This may be due to there simply not being a relationship between these variables among Black Americans, but it may also be due to a lack of statistical power.

4 Conclusion

While it has long been hypothesized that beliefs about the ethnic distribution of welfare recipients influences support for welfare, our study is the first to test this hypothesis using

¹⁷This data is only from experiment 2, as we did not ask this question in experiment 1.

the tools of randomized inference. We find strong evidence, across two separate experiments, that increasing beliefs about the share who are Black reduces support for welfare among White survey respondents. Moreover, we find that priming individuals to think about race can also dampen their support for redistribution.

Our study does not directly address how these findings could affect welfare policy. Nonetheless, theoretical work suggests that, under political competition, the preferences of voters can have a substantial impact on equilibrium policies related to redistribution ([Lee et al. \(2006\)](#)). Based on our findings, it is therefore possible that US welfare spending may be lower than it would be if attitudes toward welfare were not affected by beliefs about the racial composition of welfare recipients.

References

- Alesina, A., Ferroni, M. F., and Stantcheva, S. (2021). Perceptions of racial gaps, their causes, and ways to reduce them. Technical report, National Bureau of Economic Research.
- Alesina, A. and Glaeser, E. L. (2004). *Fighting poverty in the US and Europe: A world of difference*. Oxford University Press.
- Alesina, A., Miano, A., and Stantcheva, S. (2018). Immigration and redistribution. Technical report, National Bureau of Economic Research.
- Alesina, A. and Stantcheva, S. (2020). Diversity, immigration, and redistribution. In *AEA Papers and Proceedings*, volume 110, pages 329–34.
- Alesina, A. F., Glaeser, E. L., and Sacerdote, B. (2001). Why doesn't the us have a european-style welfare system?
- Altonji, J. G. and Doraszelski, U. (2005). The role of permanent income and demographics in black/white differences in wealth. *Journal of Human Resources*, 40(1):1–30.
- Ashok, V., Kuziemko, I., and Washington, E. (2015). Support for redistribution in an age of rising inequality: New stylized facts and some tentative explanations. Technical report, National Bureau of Economic Research.
- Augenblick, N. and Rabin, M. (2021). Belief movement, uncertainty reduction, and rational updating. *Quarterly Journal of Economics*, 136(2):933–985.
- Barsky, R., Bound, J., Charles, K. K., and Lupton, J. P. (2002). Accounting for the black–white wealth gap: a nonparametric approach. *Journal of the American statistical Association*, 97(459):663–673.
- Bobo, L. and Kluegel, J. R. (1993). Opposition to race-targeting: self-interest, stratification ideology, or racial attitudes? *American Sociological Review*, pages 443–464.
- Brodish, A. B., Brazy, P. C., and Devine, P. G. (2008). More eyes on the prize: Variability in white americans' perceptions of progress toward racial equality. *Personality and Social Psychology Bulletin*, 34(4):513–527.
- Bursztyn, L., Egorov, G., and Fiorin, S. (2020a). From extreme to mainstream: The erosion of social norms. *American Economic Review*, 110(11):3522–48.

- Bursztyn, L., Haaland, I. K., Rao, A., and Roth, C. P. (2020b). Disguising prejudice: Popular rationales as excuses for intolerant expression. Technical report, National Bureau of Economic Research.
- Callaghan, B., Harouni, L., Dupree, C. H., Kraus, M. W., and Richeson, J. A. (2021). Testing the efficacy of three informational interventions for reducing misperceptions of the black–white wealth gap. *Proceedings of the National Academy of Sciences*, 118(38).
- Charles, K. K. and Hurst, E. (2002). The transition to home ownership and the black-white wealth gap. *Review of Economics and Statistics*, 84(2):281–297.
- Chiteji, N. S. and Stafford, F. P. (1999). Portfolio choices of parents and their children as young adults: Asset accumulation by african-american families. *American Economic Review*, 89(2):377–380.
- Cook, F. L. and Barrett, E. J. (1992). *Support for the American welfare state: The views of Congress and the public*. Columbia University Press.
- Crowne, D. P. and Marlowe, D. (1960). A new scale of social desirability independent of psychopathology. *Journal of Consulting Psychology*, 24(4):349.
- Dahlberg, M., Edmark, K., and Lundqvist, H. (2012). Ethnic diversity and preferences for redistribution. *Journal of Political Economy*, 120(1):41–76.
- Derenoncourt, E., Kim, C. H., Kuhn, M., and Schularick, M. (2021). The racial wealth gap, 1860-2020. *Manuscript, Princeton University and University of Bonn*.
- Dhar, D., Jain, T., and Jayachandran, S. (2018). Reshaping adolescents’ gender attitudes: Evidence from a school-based experiment in india. Technical report, National Bureau of Economic Research.
- Eibach, R. P. and Keegan, T. (2006). Free at last? social dominance, loss aversion, and white and black americans’ differing assessments of racial progress. *Journal of Personality and Social Psychology*, 90(3):453.
- Flory, J. A., Leibbrandt, A., and List, J. A. (2015). Do competitive workplaces deter female workers? a large-scale natural field experiment on job entry decisions. *Review of Economic Studies*, 82(1):122–155.
- Fong, C. (2001). Social preferences, self-interest, and the demand for redistribution. *Journal of Public Economics*, 82(2):225–246.

- Fong, C. M. and Luttmer, E. F. (2009a). What determines giving to hurricane katrina victims? experimental evidence on racial group loyalty. *American Economic Journal: Applied Economics*, 1(2):64–87.
- Fong, C. M. and Luttmer, E. F. (2009b). What determines giving to hurricane katrina victims? experimental evidence on racial group loyalty. *American Economic Journal: Applied Economics*, 1(2):64–87.
- Fong, C. M. and Luttmer, E. F. (2011). Do fairness and race matter in generosity? evidence from a nationally representative charity experiment. *Journal of Public Economics*, 95(5-6):372–394.
- Gilens, M. (1995). Racial attitudes and opposition to welfare. *Journal of Politics*, 57(4):994–1014.
- Gilens, M. (1996). “race coding” and white opposition to welfare. *American Political Science Review*, 90(3):593–604.
- Gilens, M. (2009). *Why Americans hate welfare: Race, media, and the politics of antipoverty policy*. University of Chicago Press.
- Gittleman, M. and Wolff, E. N. (2004). Racial differences in patterns of wealth accumulation. *Journal of Human Resources*, 39(1):193–227.
- Gross, K. and Wronski, J. (2021). Helping the homeless: The role of empathy, race and deservingness in motivating policy support and charitable giving. *Political Behavior*, 43(2):585–613.
- Haaland, I. and Roth, C. (2021). Beliefs about racial discrimination and support for pro-black policies. *Review of Economics and Statistics*, pages 1–38.
- Haaland, I., Roth, C., and Wohlfart, J. (2020). Designing information provision experiments.
- Harell, A., Soroka, S., and Iyengar, S. (2016). Race, prejudice and attitudes toward redistribution: A comparative experimental approach. *European Journal of Political Research*, 55(4):723–744.
- Henry, P., Reyna, C., and Weiner, B. (2004). Hate welfare but help the poor: how the attributional content of stereotypes explains the paradox of reactions to the destitute in america 1. *Journal of Applied Social Psychology*, 34(1):34–58.

- Kraus, M. W., Onyeador, I. N., Daumeyer, N. M., Rucker, J. M., and Richeson, J. A. (2019). The misperception of racial economic inequality. *Perspectives on Psychological Science*, 14(6):899–921.
- Kraus, M. W., Rucker, J. M., and Richeson, J. A. (2017). Americans misperceive racial economic equality. *Proceedings of the National Academy of Sciences*, 114(39):10324–10331.
- Lee, W., Roemer, J., and Van der Straeten, K. (2006). Racism, xenophobia, and redistribution. *Journal of the European Economic Association*, 4(2-3):446–454.
- Lee, W. and Roemer, J. E. (2006). Racism and redistribution in the united states: A solution to the problem of american exceptionalism. *Journal of Public Economics*, 90(6-7):1027–1052.
- List, J. A. (2004). The nature and extent of discrimination in the marketplace: Evidence from the field. *Quarterly Journal of Economics*, 119(1):49–89.
- Luttmer, E. F. (2001). Group loyalty and the taste for redistribution. *Journal of Political Economy*, 109(3):500–528.
- Onyeador, I. N., Daumeyer, N. M., Rucker, J. M., Duker, A., Kraus, M. W., and Richeson, J. A. (2021). Disrupting beliefs in racial progress: Reminders of persistent racism alter perceptions of past, but not current, racial economic equality. *Personality and Social Psychology Bulletin*, 47(5):753–765.
- O’Brien, R. L. (2017). Redistribution and the new fiscal sociology: Race and the progressivity of state and local taxes. *American Journal of Sociology*, 122(4):1015–1049.
- Peer, E., Brandimarte, L., Samat, S., and Acquisti, A. (2017). Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163.
- Phillips, L. T. and Lowery, B. S. (2015). The hard-knock life? whites claim hardships in response to racial inequity. *Journal of Experimental Social Psychology*, 61:12–18.
- Quadagno, J. S. et al. (1994). *The color of welfare: How racism undermined the war on poverty*. Oxford University Press.
- Ribar, D. C. and Wilhelm, M. O. (1999). The demand for welfare generosity. *Review of Economics and Statistics*, 81(1):96–108.
- Rothstein, R. (2017). *The color of law: A forgotten history of how our government segregated America*. Liveright Publishing.

Wetts, R. and Willer, R. (2018). Privilege on the precipice: Perceived racial status threats lead white americans to oppose welfare programs. *Social Forces*, 97(2):793–822.

Appendices

A Tables and figures

Table A1: Balance table (experiment 1)

	[1]	[2]	[3]	[1] vs [2]	[1] vs [3]	[2] vs [3]	Joint test
Female	0.584	0.597	0.597	0.433	0.409	0.965	0.649
Age	34.5	34.4	33.6	0.764	0.023	0.049	0.050
Moderate	0.259	0.267	0.275	0.546	0.251	0.584	0.517
Conservative	0.153	0.168	0.145	0.200	0.478	0.046	0.128
Ever on welfare	0.564	0.542	0.554	0.181	0.540	0.469	0.407
College educated	0.593	0.589	0.587	0.793	0.699	0.900	0.925
Voted	0.840	0.845	0.846	0.702	0.636	0.927	0.881
<i>n</i>	1926	1943	1924				

Table A2: Balance table (experiment 2)

	[1]	[2]	[3]	[1] vs [2]	[1] vs [3]	[2] vs [3]	Joint test
Female	0.680	0.644	0.663	0.047	0.357	0.288	0.139
Age	35.4	35.9	35.3	0.365	0.949	0.326	0.546
Moderate	0.321	0.336	0.356	0.402	0.052	0.271	0.150
Conservative	0.321	0.305	0.291	0.384	0.094	0.423	0.246
Ever on welfare	0.423	0.418	0.436	0.812	0.513	0.373	0.652
College educated	0.548	0.563	0.553	0.438	0.779	0.620	0.734
Voted	0.842	0.829	0.826	0.369	0.259	0.818	0.494
<i>n</i>	1331	1324	1327				

Table A3: Beliefs by sub-group

Subgroup	% Black	% White	% Latino
White respondents	37.52 (14.54)	38.42 (15.87)	26.72 (15.25)
Black respondents	35.35 (15.26)	38.92 (17.32)	28.63 (15.39)
Female respondents	38.16 (14.63)	37.86 (15.76)	27.58 (15.25)
Male respondents	35.39 (14.65)	39.60 (16.72)	26.26 (15.34)
Conservative respondents	40.57 (15.71)	35.79 (15.40)	28.22 (16.65)
Moderate respondents	38.38 (14.59)	36.55 (15.49)	27.74 (15.38)
Liberal respondents	34.80 (13.89)	40.92 (16.52)	26.17 (14.55)
With a college degree	36.03 (14.83)	40.34 (16.41)	26.64 (15.58)
Without a college degree	38.58 (14.40)	36.05 (15.46)	27.68 (14.88)
Have ever been on welfare	36.00 (14.44)	39.38 (16.21)	28.19 (16.45)
Have never been on welfare	38.23 (14.88)	37.65 (16.05)	25.97 (13.96)
Voted	36.62 (14.63)	39.10 (16.23)	26.80 (15.30)
Did not vote	39.68 (14.80)	35.48 (15.40)	28.55 (15.16)

Table A4: Predictors of beliefs

Variables	Beliefs about % who are Black
Female	3.002*** (0.298)
Age	-0.0867*** (0.0116)
Moderate	3.391*** (0.342)
Conservative	6.056*** (0.400)
Ever on welfare	-2.043*** (0.296)
College	-2.137*** (0.305)
Voted	-1.696*** (0.421)
Black	-1.822*** (0.391)
Constant	39.96*** (0.590)
<i>n</i>	9,775
<i>R</i> ²	0.058

Table A5: Bayesian updating

	log(% welfare recipients that are black)
log(% that are Black)	0.242*** (0.00845)
log(% Black Americans on welfare)	0.623*** (0.00711)
log(% Americans on welfare)	-0.427*** (0.00842)
<i>n</i>	9,765
<i>R</i> ²	0.984

Table A6: Treatment effects on beliefs

Variables	% who are Black (exp 1)	% who are Black (exp 2)
High estimate	-4.235*** (0.527)	11.34*** (0.438)
Low estimate	-8.403*** (0.556)	-25.51*** (0.489)
Constant	35.16*** (0.401)	39.48*** (0.394)
<i>n</i>	3,948	3,981
<i>R</i> ²	0.059	0.660

Table A7: Confidence in beliefs

Variables	Confident (exp. 1)	Confident (exp. 2)	Confident (exp. 1)	Confident (exp. 2)
High estimate	-0.0334* (0.0193)	0.0278 (0.0188)	0.0529*** (0.0191)	0.0655*** (0.0185)
Low estimate			0.0864*** (0.0192)	0.0377** (0.0184)
Constant	0.459*** (0.0137)	0.361*** (0.0132)	0.373*** (0.0134)	0.323*** (0.0128)
<i>n</i>	2,646	2,651	3,948	3,981
<i>R</i> ²	0.001	0.001	0.005	0.003

Table A8: Pooled estimates

	Pro welfare		Anti welfare	
High estimate	-0.0387**		0.00118	
	(0.0164)		(0.0108)	
Experiment 1	0.0160		0.0298**	
	(0.0176)		(0.0116)	
High estimate \times exp 1	0.00556		0.0307*	
	(0.0243)		(0.0166)	
Female	0.0235*	0.0226*	-0.0211**	-0.0268***
	(0.0125)	(0.0124)	(0.00870)	(0.00876)
Age	-0.00107**	-0.00103**	-0.000318	-0.000242
	(0.000489)	(0.000483)	(0.000340)	(0.000338)
Moderate	-0.163***	-0.165***	0.0389***	0.0355***
	(0.0158)	(0.0155)	(0.0103)	(0.00993)
Conservative	-0.304***	-0.305***	0.0781***	0.0814***
	(0.0162)	(0.0141)	(0.0131)	(0.0117)
Ever on welfare	0.0401***	0.0409***	-0.0160*	-0.0100
	(0.0125)	(0.0125)	(0.00855)	(0.00853)
College	0.00716	0.00708	0.0150*	0.0180**
	(0.0128)	(0.0128)	(0.00894)	(0.00893)
Right wing news	-0.00186		0.0300***	
	(0.0144)		(0.0113)	
Voted	0.0109	0.0112	0.0260**	0.0286**
	(0.0179)	(0.0179)	(0.0118)	(0.0119)
Belief about % black		-0.00120***		-4.37e-05
		(0.000439)		(0.000292)
Constant	0.433***	0.459***	0.0424**	0.0694***
	(0.0276)	(0.0285)	(0.0185)	(0.0193)
n	5,297	5,297	5,297	5,297
R^2	0.085	0.083	0.024	0.015

Table A9: Robustness checks (experiment 1)

Robustness check	Coefficient	Standard error	<i>n</i>
Drop if time taken <6min	0.00629**	0.00316	2,571
" 7min	0.00580*	0.00318	2,422
" 8min	0.00611*	0.00323	2,235
" 9min	0.00512	0.00326	1,996
" 10min	0.00332	0.00341	1,684
Drop if treat time <5sec	0.00649**	0.00316	2,554
" 10sec	0.00459	0.00304	1,994
" 15sec	0.00156	0.00379	1,300
" 20sec	-0.000930	0.00423	823
Drop if failed quiz	0.00413	0.0088	1,209
Drop if belief = 0/100	0.00716**	0.00315	2,646
Drop if % win lottery <1%	0.00703**	0.00337	2,457
" 2%	0.00990**	0.00438	1,810
" 3%	0.00975**	0.00461	1,591
" 4%	0.00906*	0.0049	1,424
" 5%	0.00875*	0.00489	1,313
Drop if understood purpose	0.00677*	0.00408	1,293
Re-weighted analysis	0.00457	0.0034	2,646
Control for confidence	0.00739**	0.00315	2,646
Instrument for confidence	0.0263	0.0465	3,948
Use implied beliefs	0.0300	0.0207	2,645
Include control group	0.00178	0.0015	3,948

Table A10: Robustness checks (experiment 2)

Robustness check	Coefficient	Standard error	<i>n</i>
Drop if time taken <6min	-0.000972**	0.000472	2,626
” 7min	-0.000931*	0.000481	2,556
” 8min	-0.000864*	0.000500	2,410
” 9min	-0.00111**	0.000533	2,196
” 10min	-0.00121**	0.000574	1,943
Drop if treat time <5sec	-0.000979**	0.000469	2,651
” 10sec	-0.000979**	0.000469	2,651
” 15sec	-0.000979**	0.000469	2,651
” 20sec	-0.000979**	0.000469	2,651
Drop if failed quiz	-0.000360	0.000844	800
Drop if belief = 0/100	-0.000979**	0.000469	2,651
Drop soc. des. bias	-0.00127**	0.000512	2,153
Drop if understood purpose	-0.000928*	0.000527	1,981
Re-weighted analysis	-0.00162***	0.000547	2,651
Control for confidence	-0.000951**	0.000469	2,651
Instrument for confidence	-0.000835*	0.000481	3,981
Use implied beliefs	-0.00140**	0.000674	2,650
Include control group	-0.00907**	0.0004603	3,981

Table A11: Association between beliefs and welfare support for White respondents

Variables	Pro-welfare donation	Anti-welfare donation	Pro-welfare donation	Anti-welfare donation
Belief % Black (post-treatment)	-0.000679** (0.000293)	8.83e-05 (0.000205)		
Belief % Black (pre-treatment)			-0.00106*** (0.000346)	0.000267 (0.000250)
Female	0.0115 (0.0102)	-0.0253*** (0.00712)	0.0139 (0.0103)	-0.0261*** (0.00713)
Age	-0.00136*** (0.000390)	-0.000239 (0.000278)	-0.00142*** (0.000391)	-0.000219 (0.000280)
Moderate	-0.180*** (0.0126)	0.0398*** (0.00805)	-0.178*** (0.0126)	0.0390*** (0.00810)
Conservative	-0.301*** (0.0116)	0.0851*** (0.00952)	-0.297*** (0.0118)	0.0838*** (0.00970)
Ever on welfare	0.0435*** (0.0102)	-0.00610 (0.00698)	0.0416*** (0.0103)	-0.00555 (0.00694)
College	0.0196* (0.0105)	0.0202*** (0.00721)	0.0184* (0.0105)	0.0207*** (0.00720)
Voted	0.0174 (0.0147)	0.0206** (0.00979)	0.0168 (0.0147)	0.0208** (0.00978)
Constant	0.454*** (0.0224)	0.0647*** (0.0152)	0.473*** (0.0243)	0.0572*** (0.0166)
<i>n</i>	7,929	7,929	7,929	7,929
<i>R</i> ²	0.085	0.016	0.086	0.016

Table A12: Predictors of welfare support

Variables	Pro-welfare donation	Anti-welfare donation
Female	-0.00215 (0.00944)	-0.0274*** (0.00648)
Age	-0.00152*** (0.000364)	-0.000344 (0.000249)
Moderate	-0.161*** (0.0112)	0.0392*** (0.00722)
Conservative	-0.279*** (0.0110)	0.0832*** (0.00893)
Ever on welfare	0.0445*** (0.00941)	-0.00343 (0.00631)
College	0.0155 (0.00966)	0.0182*** (0.00647)
Vote	0.0140 (0.0130)	0.0192** (0.00855)
Black	0.0142 (0.0128)	0.0121 (0.00784)
Constant	0.440*** (0.0189)	0.0744*** (0.0124)
<i>n</i>	9,775	9,775
<i>R</i> ²	0.069	0.015

Table A13: ITT and LATE for Black respondents

Variables	Pro-welfare donation	Pro-welfare donation	Anti-welfare donation	Anti-welfare donation
High estimate	-0.00567 (0.0276)		-0.00236 (0.0175)	
Beliefs about the % who are Black		-0.00151 (0.00736)		-0.000628 (0.00466)
Constant	0.372*** (0.0195)	0.414* (0.218)	0.105*** (0.0124)	0.123 (0.138)
<i>n</i>	1,221	1,221	1,221	1,221
<i>R</i> ²	0.000		0.000	