



# Are Americans less likely to reply to emails from Black people relative to White people?

Ray Block Jr.<sup>a,1</sup>, Charles Crabtree<sup>b,1</sup>, John B. Holbein<sup>c,1,2</sup>, and J. Quin Monson<sup>d,1</sup>

<sup>a</sup>Department of Political Science, Pennsylvania State University, University Park, PA 16802; <sup>b</sup>Department of Government, Dartmouth College, Hanover, NH 03755; <sup>c</sup>Frank Batten School of Leadership & Public Policy, University of Virginia, Charlottesville, VA 22904; and <sup>d</sup>Department of Political Science, Brigham Young University, Provo, UT 84602

Edited by Margaret Levi, Department of Political Science, Stanford University, Stanford, CA; received June 3, 2021; accepted October 29, 2021

**In this article, we present the results from a large-scale field experiment designed to measure racial discrimination among the American public. We conducted an audit study on the general public—sending correspondence to 250,000 citizens randomly drawn from public voter registration lists. Our within-subjects experimental design tested the public’s responsiveness to electronically delivered requests to volunteer their time to help with completing a simple task—taking a survey. We randomized whether the request came from either an ostensibly Black or an ostensibly White sender. We provide evidence that in electronic interactions, on average, the public is less likely to respond to emails from people they believe to be Black (rather than White). Our results give us a snapshot of a subtle form of racial bias that is systemic in the United States. What we term everyday or “paper cut” discrimination is exhibited by all racial/ethnic subgroups—outside of Black people themselves—and is present in all geographic regions in the United States. We benchmark paper cut discrimination among the public to estimates of discrimination among various groups of social elites. We show that discrimination among the public occurs more frequently than discrimination observed among elected officials and discrimination in higher education and the medical sector but simultaneously, less frequently than discrimination in housing and employment contexts. Our results provide a window into the discrimination that Black people in the United States face in day-to-day interactions with their fellow citizens.**

racial bias | audit study | discrimination against African Americans

classmates, coworkers, and kinfolk to the many types of electronic communication that have become so frequent in modern society. Although opportunities for overt racial hostility are all too common, these are comparatively rarer when benchmarked against the many ways that people spend their time in their day-to-day lives. Furthermore, while overt instances of racial discrimination have been widely documented and studied, much less understood is the extent, nature, and origins of subtler, but pernicious, “everyday racism” (33). These common, seemingly benign, social interactions that are (potentially) a breeding ground for what we call “paper cut” discrimination happen countless times each day. The people doing the cutting may not always be aware of their actions, and the ones being cut often struggle to identify and/or respond to these slights (34). While these indignities may seem minor, their cumulative impact is likely substantial; indeed, as scholars of racial microaggressions remind us, seemingly small acts of discrimination are important to consider (35).

To what extent do Americans of various backgrounds engage in paper cut discrimination against Black people? To be clear, what we (and others in this literature) mean when we use the term racial discrimination is the differential treatment of individuals who are all else equal, apart from their racial identity. Although in common parlance, “discrimination” is sometimes used to impute motivations or intentions that are internal to the subjects being studied, that is not how we (or others in this literature) use this term. As prominent social scientists who led the charge in

In his influential book *The Souls of Black Folk: Essays and Reflections*, W. E. B. Du Bois notes that the “problem of the twentieth century is the problem of the color line” (1), and he characterizes this color line as “the question of how far differences of race . . . will be made hereafter the basis of denying to over half the world the right of sharing to their utmost ability the opportunities and privileges of modern civilization” (2). Unfortunately, Du Bois’ words remain relevant in the twenty-first century, for the United States (among many other countries) continues to grapple with this “question of racial differences.” In fact, one of America’s enduring legacies is the many active and antagonistic acts of racial bias visited upon its Black citizens. Violence and overt hostility against African Americans have been, and continue to be, integral parts of the story of the United States. From historical patterns of slavery (3–6), Jim Crow (7, 8), and de facto and de jure segregation (9) to modern manifestations of police brutality (10, 11), racialized patterns of incarceration (12), unequal treatment in the labor force (13–15), and racial inequities in public services (7, 16–19), African Americans are acutely experienced with myriad forms of racial bias, each of which cuts deeply and violently into their lives.

Although volumes of research consider these deep cuts (19–32), prior work has struggled to capture more subtle types of discrimination against Black people. These subtler forms of discrimination may manifest themselves frequently in the simple, oft-observed social interactions that make up our everyday lives—from exchanges with one’s neighbors, churchgoers,

## Significance

Although previous attempts have been made to measure everyday discrimination against African Americans, these approaches have been constrained by distinct methodological challenges. We present the results from an audit or correspondence study of a large-scale, nationally representative pool of the American public. We provide evidence that in simple day-to-day interactions, such as sending and responding to emails, the public discriminates against Black people. This discrimination is present among all racial/ethnic groups (aside from among Black people) and all areas of the country. Our results provide a window into the discrimination that Black people in the United States face in day-to-day interactions with their fellow citizens.

Author contributions: R.B., C.C., J.B.H., and J.Q.M. designed research, performed research, contributed new reagents/analytic tools, analyzed data, and wrote the paper. The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

<sup>1</sup>R.B., C.C., J.B.H., and J.Q.M. contributed equally to this work.

<sup>2</sup>To whom correspondence may be addressed. Email: holbein@virginia.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2110347118/-DCSupplemental>.

Published December 20, 2021.

using audit studies on social elites, Pager and Shepherd (36, p. 182) note that

[r]acial discrimination refers to unequal treatment of persons or groups on the basis of their race or ethnicity . . . A key feature of any definition of discrimination is its focus on behavior. Discrimination is distinct from racial prejudice (attitudes), racial stereotypes (beliefs), and racism (ideologies) that may also be associated with racial disadvantage. Discrimination may be motivated by prejudice, stereotypes, or racism, but the definition of discrimination does not presume [an] underlying cause.

According to Pager and Shepherd (36), discrimination does not subsume a singular motive. Indeed, discrimination may be driven by a bundle of mechanisms, including “taste-based” or “statistical” considerations (37), pro-White and/or anti-Black attitudes (38), or something else entirely. Regardless of the exact mechanisms, the results of racial discrimination are the same—one disadvantaged group gets inferior treatment.

Although previous attempts have been made to measure racial discrimination among the public, this work has faced distinct and difficult methodological challenges. These challenges are—in large part—due to the inherent complications of observing and measuring sensitive attitudes/behaviors (39). To our knowledge, no previous work has experimentally tested for racial discrimination in the real-world behavior of a large, nationally representative sample of the American public. Instead, previous work has focused on measuring discrimination among social elites and people in certain occupations—from elected officials to those who oversee the employment and housing sector to medical professionals and to those involved in educating young people (13, 14, 16–18, 40–43).

In this paper, we address this opening in the literature. Specifically, our approach to studying everyday racial discrimination uses a standardized request for help (in this case, an email request to volunteer one’s time by taking a short survey) in which subjects are randomly assigned to one of two treatment arms—receiving the request from either a Black requester or a White requester. This request has close parallels to the information/service requests commonly found in audit studies of elites and as we validate below, elicits similar behavior among this group. Our specific request involves an invitation to volunteer one’s time by completing a brief survey. After sending the email, we then observe whether there are differential rates of helping behavior (i.e., response and completion rates of the survey) depending on the race of the sender.

We describe this task in greater detail below; however, we (briefly) note that this approach circumvents several issues that have limited prior research. First and foremost, we measure discrimination by directly examining citizens’ racial actions as opposed to measuring indirect proxies of their racial attitudes. We do so by providing a behavioral measure of racial discrimination that is taken in a real-world setting and not an artificial laboratory. Further, our approach avoids issues of Hawthorne effects potentially present in laboratory or survey environments because the subjects in our study do not know they are being studied. Moreover, our task is not confounded by other potential sources of bias—people in our research design are given the same exact task, with no differences across the task other than the randomly assigned racial manipulation. Our study accomplishes all of these goals with a large and nationally representative sample of the public, which allows us to make inferences not only at the national level but also, in each of the 50 US states individually. Finally, our work helps benchmark discrimination among the public to discrimination observed among various social elites. We do so by 1) providing a direct comparison of our experimental effects with those observed in an additional audit study that we

run on an important group of social elites—elected officials—and 2) drawing from a meta-analysis that outlines the extent of discrimination shown in audit studies among various other social elites—including those elites who oversee the employment, housing, health, and education sectors of our modern society. This allows us to see whether the public mirrors or contrasts the behavior of social elites among whom previous research has focused.

## Experimental Design and Methods

To measure citizens’ levels of everyday or paper cut discrimination, we conducted an audit study on a pool of randomly drawn citizens (pairing it with an audit study of elected officials). (As we outline in *SI Appendix*, sections 1.5 and 1.6, our work is comparable with, but ultimately distinct from, audits of employers, studies of the race of interviewer effect on response rates, and studies that use the “Lost Letter” technique.) Our audit study—like previous studies of public officials (16–18, 40–42, 44) and other social elites (13, 14, 40, 43)—consisted of us reaching out to subjects with a simple volunteer-based request for them to provide information/service/help. To ensure that our request made sense to both the public and elected officials, our experiment was couched within a request to volunteer to take a survey on contemporary political issues. Although at first blush, this design may seem different from previous audit studies, our request closely parallels the informational/service requests commonly found in audit studies of elites. We elaborate on how our request is similar to those used in other audit studies in *SI Appendix*, sections 1.5, 2, and 6. However, we note here that we empirically validate that this request elicits similar behavior to the informational/service requests used in previous audit studies of social elites. At its core, our invitation consists of a simple request to volunteer one’s time and effort to help in this specific online domain.

Our invitation came from a nondescript survey firm. Unlike some audit studies, our request was not fictional (45). We were fielding an actual survey. This allowed a group of scholars to collect information on where a large sample of public and elected officials stood on a number of contemporary social and political issues. [As we outline further in *SI Appendix*, section 7, this intentional design feature further adds to the value of our study (46).] We sent the emails using Qualtrics. We randomized the name of the sender of the survey invitation to be either putatively White or putatively Black (47). (We also randomized the order in which we sent the emails.) As we discuss in *SI Appendix*, we used names that are predominantly White or Black in government records and that have been shown to also be perceived as predominantly Black/White by the public (48, 49). All other information about the requester—including socioeconomic and gender status—was held constant.

Our public sample consisted of 250,000 randomly chosen individuals from a nationwide voter registration list compiled and maintained by the data and analytics firm DT Client Services, LLC (commonly known as “The Data Trust”). For those unfamiliar, in the United States voter registration lists are public record. These files are collated together by firms like the one we used in this study. Voter registration lists include the vast majority of American citizens and for this reason, have been used to study a wide variety of social phenomena (50–52). Depending on the data source one uses, voter registration lists contain  $\approx 80\%$  of the adult population. Although they do not cover everyone, voter lists have one of the largest coverage rates among readily available datasets. For this reason, they are frequently used by survey firms to measure the public’s attitudes on a host of social issues. Research has shown that samples from voter lists produce estimates that are indistinguishable from other sampling strategies (e.g., random digit dialing) (53). To allow us to make comparisons within states, using an ex ante power analysis we

targeted 5,000 as our desired number of randomly selected individuals in all 50 states. To compile the contact information for 5,000 citizens per state, we oversampled citizens from the voter lists proportional to the email coverage available in these lists. We then purchased emails from The Data Trust. Their email list is drawn from a variety of commercial and public data sources. For cost purposes and to align with our desired sample size, The Data Trust stopped matching to emails once we hit our 5,000-individuals target in each state. (We chose an equal number of individuals to equally power our treatment effects across states. However, in *SI Appendix*, we show that our results are robust to using state population weights [*SI Appendix*, Fig. S17].)

As a benchmark for our public estimates, we replicate our study design with elected officials. This intentional design feature (along with pairing our estimates with previous meta-analytic averages of discrimination levels among various social elites) not only allows us to get a sense of whether our effects among the public are large or are small but also, allows us to explore questions of whether previously documented levels of discrimination among elected officials (*SI Appendix*, section 1.1 has a review of this literature and our own original meta-analysis of study estimates therein) are in concordance with what their constituents themselves exhibit. (We return to the theoretical contribution of benchmarking public and elected official discrimination estimates in *SI Appendix*, sections 1.2 and 1.3.) Our elected official sample consists of all state legislators, mayors, and city councilors in the United States ( $N \approx 40,000$ ). To ensure that we did not cross-contaminate our samples, in local offices we clustered our treatments among local officials at the municipality level. Data on these political elites were scraped from publicly available sources (via <https://openstates.org/>) and collected from an email list provided to us by the American Municipal Officials Survey (54). We included all of these public officials at the state and local levels so that we may generalize our findings to a nationally representative pool of elected officials nested in the same states as the public sample. Like the public sample, this large sample also allows us to make (well-powered) estimates at the national and subnational levels. It also allows us to validate that our specific request produces similar estimates to the other requests used in the audit study literature.

We intentionally used large samples of both the public and elected officials for four core reasons. First, we wanted to ensure that our estimates are sufficiently powered to avoid errors that often arise in underpowered research designs (55). Second, we chose to use a large sample to help ensure that we have enough statistical power to identify effects given the *ex ante* likelihood of floor effects in the response rates of our elite and public surveys. Our sample sizes were chosen by a power analysis conducted at the design stage of the experiment that assumed a 1% overall response rate in line with the experiences of other researchers. In practice, the overall response rate among emails that did not bounce (due to being invalid addresses) for the public was  $\sim 2.0\%$ , and the response rate for elected officials was  $\sim 5.7\%$ . Third, we chose our large sample sizes to allow the consortium of scholars whose questions were embedded in the survey to get a sufficiently large sample; given the expected low response rates and their desired sample size, we needed to send a large number of invitations. Fourth, we chose a large sample size because of our (preregistered) desire to look for differences in response types across states and other individual characteristics. Previous research has shown that estimating treatment effect heterogeneity requires more statistical power than many researchers suspect, and thus, many subgroup tests are underpowered (56–58).

To further increase the robustness of our design and to ensure that our effects do not vary by the timing of the study, we employ a within-subjects design for both the public and elected official samples. This means that all subjects received emails to two

separate surveys. As in all within-subjects designs, the ordering of the treatments in our study is randomly assigned (59), thus removing context between the periods of study as a reason for the effects observed. The two survey invitations were separated by a roughly 3-mo period—the first wave was sent at the end of January 2020, and the second wave was sent at the end of April 2020 (i.e., both waves were conducted before the George Floyd protests that began in May and June of 2020). Combining our data across these two survey waves, our end sample consists of 500,000 citizen-wave observations for the public and just over 81,000 elected official-wave observations. Our unit of analysis with this sample is the individual wave. In our models, we include individual fixed effects as well as indicators for which treatment the individual was assigned in that wave. (In practice, these turn out to be slightly more precise than models that include fixed effects for blocks. However, as we show in *SI Appendix*, Figs. S19 and S20, both approaches yield similar conclusions.)

As in audit studies of elites, our key outcome of interest is whether the recipients responded to the email request (i.e., whether they clicked an embedded link to the first question of the survey, thus choosing to volunteer their time and reply to our request for help with providing the requested information embedded in the survey questions). We focus on whether a recipient responded as opposed to measures of response quality, like response tone, given that focusing on the latter can induce posttreatment bias (60, 61). That is, in only being measured among those who take the survey, many measures of response quality effectively break the randomization of the experimental design. These problems are unavoidable with many traditional measures of response quality. Hence, conditioning on response in an audit study is generally a bad idea (60). That said, measures of response quality can be used, with care—as long as one does not condition on response. For instance, we examine the extent to which participants not only started the survey but whether they also finished it (with nonrespondents being marked as neither starting nor finishing the survey) and find substantively similar results. Taken literally, these measure respondents' willingness to respond to and complete surveys based on whether the requester is White or Black. However, on a more basic and abstract level, this provides us with a behavioral measure of the public's decision to volunteer their time and effort (i.e., to provide time and information) to the senders of our emails.

Our project was approved by the Brigham Young University Institutional Review Board (E17512 and X19052), which waived informed consent.

## Results

Black senders received fewer responses than White senders. Overall, among the public, 1.6% in the White sender group responded; in the Black sender group, 1.4% of the public responded. Put in terms of counts, the Black sender received 3,620 responses (of 250,000), whereas the White sender received 4,007 responses (of 250,000). While the substantive significance of effects is always, to varying degrees, in the eye of the beholder (62), there are several comparison points that we can provide to give a sense for how large or small this difference in response patterns is.

First, the response patterns we observe correspond to an individual fixed effects adjusted odds ratio (OR) of 1.155 (95% CI: [1.095, 1.219]; without fixed effects OR: 1.109; 95% CI: [1.06, 1.16])—meaning that the odds that the White sender received a response were (on average) 15.5% higher than the odds of receiving a response from the Black sender. The risk ratio for the White sender coefficient is 1.11 (95% CI: [1.06, 1.16])—meaning that response was increased by a factor of (approximately) 1.11 for senders who were White vs. senders who were Black.

Another comparison we include to assess how large/small our effects are is to benchmark them to the base response rate in

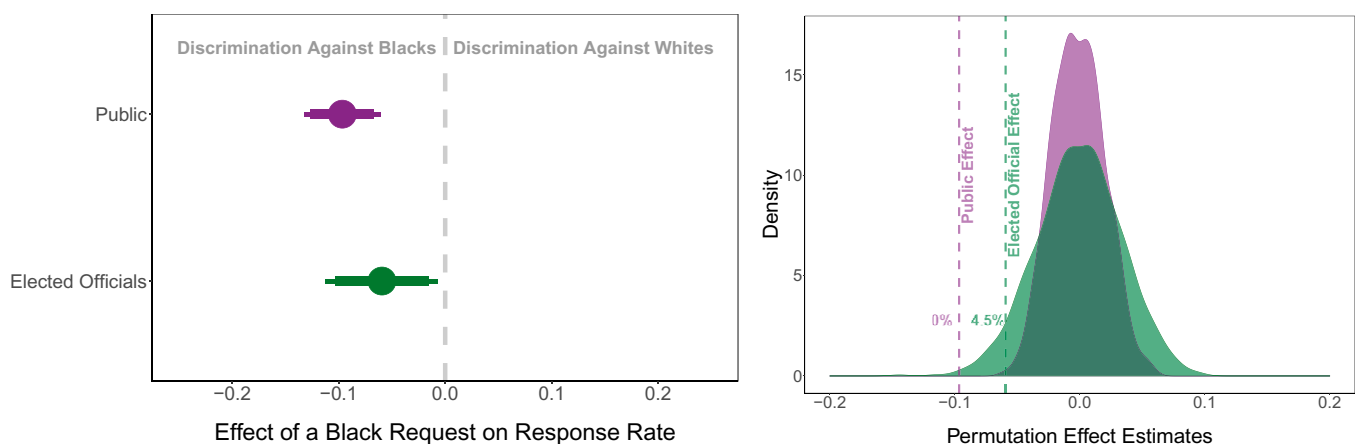
our control group. We do this because in designing our study, we anticipated that floor effects would come into play. [Comparing with base rates is a commonly used technique in economics, political science, and other disciplines (e.g., refs. 63–67).] The discrimination effect among the public is equivalent to a (on average) 9.7% discriminatory response toward African American senders. This effect is statistically significant ( $P < 0.0001$ , which is also significant at the Benjamini–Hochberg multiple comparisons adjusted critical value: 0.017) and is precisely estimated (95% CI: [−13.2%, −6.1%]).

A third way of getting a sense how large/small effects are is to benchmark them to an outside group of substantive interest. Given that 1) many previous studies have studied whether politicians discriminate against Black people and 2) the public that we are studying includes the constituents who vote for/against these elected officials, we can usefully benchmark our effects to those in this literature. As can be seen in Fig. 1, we are able to replicate the finding of other audit studies of elites that elected officials discriminate on the basis of race in simple requests for help; 4.2% of elected officials in the control (White) group responded, whereas again, only 3.9% of elected officials in the treatment (Black) group responded. This produces an (individual fixed effect adjusted) OR of 1.103 (95% CI: [1.011, 1.203]) and a risk ratio of 1.07 (95% CI: [1.00, 1.14]) for White senders vs. Black senders. Relative to the base rate, elected officials are 6% less likely to respond to a Black sender than a White sender. This effect is statistically significant ( $P = 0.027$ , which is significant at the Benjamini–Hochberg multiple comparisons adjusted critical value: 0.033) and is precisely estimated (95% CI: [−11.2%, −0.7%]). It is also similar in size to previous race-based audit studies of elected officials; for instance, the average effect in our meta-analysis of previous studies is also ~6% of the base rate (*SI Appendix, Fig. S1*). This suggests that our request—volunteering to take a survey—elicits a similar type of behavior as previous informational/service requests used in the literature, such as requests for help registering to vote and informational requests related to accessing various public services. (More on the parallels between our request and those in previous studies is in *SI Appendix, sections 1.5, 2, and 6.*) In short, our replication experiment corroborates that elected officials discriminate against Black people in simple informational/help requests like the one we use in our study.

It is important to note, however, that the public effect that we document is substantively larger than the effect we observe for elected officials. This suggests the public actually discriminates more than their political representatives. This would imply that elected officials' discrimination against Black people is a moderation of what their constituents themselves do. However, despite our pooled sample being large (just under 600,000), the difference between the public and elected officials is not statistically distinct ( $P = 0.353$ , a value that is not significant at the Benjamini–Hochberg multiple comparisons adjusted critical value: 0.05). Using equivalence testing, we can rule out modest effects (95% CI: [−15.1%, 5.3%]). Based on these CIs, we can conclude, at minimum, that elected officials do not discriminate (substantially) more than the public. However, we cannot rule out that the public and elected officials discriminate against Black people at the same level. Put differently, the effects we document among the American public are just as large, if not larger, than the effects documented in audit studies of elites.

More generally, we can compare our results with those produced by audit studies looking at racial discrimination among social elites in the employment, health, housing, and education sectors. Recall that the White/Black risk ratio for our study is 1.11 for members of the public. Based on a recent meta-analysis of 74 different audit studies conducted on social elites involving White and African American identities (68), discrimination among the public falls somewhere in the middle of previous estimates. According to that work, Black people experience the most discrimination in the labor (1.27 risk ratio) and housing markets (1.17). On the other hand, however, Black people experience less discrimination in public services, higher education, and medical studies (1.0 to 1.1 risk ratios). Viewed in this light, our findings suggest that paper cut discrimination among the public falls squarely in between the types of other discrimination faced by African Americans in interactions with social elites.

The effects we document, however, need to be placed within the broader context of Black people's interactions in America. On the one hand, the paper cut discrimination we observe here is not a one-off experience but rather, a repeated occurrence, as many people engage in electronic interactions with unknown individuals over the internet. These multiple small, seemingly negligible doses of social interactions add up to many opportunities for paper cut discrimination to occur. On the other hand, we



**Fig. 1.** Discrimination against Black people compared with White people by the public and elected representatives. *Left* displays the effect of an African American sender (vs. a White sender) on response rates of the public and of their elected officials. Effects listed here are within subjects and as such, include individual fixed effects. Effects are scaled relative to the mean response rate in the control group (i.e., they are in percentage of the base rate units).  $N$  (public) = 500,000;  $N$  (elected officials) = 81,024. The distributions in *Right* show results from permutation tests that randomly shuffle the data 1,000 times and estimate a treatment effect for each random draw. For the sake of computation time, individual fixed effects are omitted in the permutations. The reference lines show the observed effects, with labels for the number of permutation draws as extreme also labeled. In both panels, green indicates elected officials, and purple indicates the public. The public discriminates against Black people, and although discrimination among the public may be slightly larger, elected officials' discrimination is (statistically) a mirror image of the public's discrimination.

do not know if our results generalize to contexts when Americans receive emails from individuals they know. More than that, the consequences of this type of discrimination are potentially less profound than those experienced in other contexts, such as the housing and rental markets.


Two other points should be made, however, about effect sizes. While our design can provide evidence of discrimination, we cannot assess whether this discrimination is driven by pro-White preferences or anti-Black biases. We also cannot assess potential mechanisms that drive this discrimination. For example, we do not know if a request for something other than a political survey could manifest in effects that are either bigger or smaller. For example, maybe the effects we observe are driven by respondents who, because they assume a Black sender has different politics than they do, feel uncomfortable revealing their views to somebody who may disagree with them (a common concern among survey researchers). Or, maybe respondents are more likely to exhibit discrimination against Black senders when they are strangers, as our fictional identities would have likely been perceived. On the other hand, perhaps the official-sounding nature of this email made people more likely to respond evenly than an email that is of a different nature.

With these caveats in mind, our results suggest that the public exhibits discrimination against African Americans compared with White people in this small, seemingly negligible interaction with their fellow citizens. This effect is robust to a host of robustness checks that we include in *SI Appendix*.

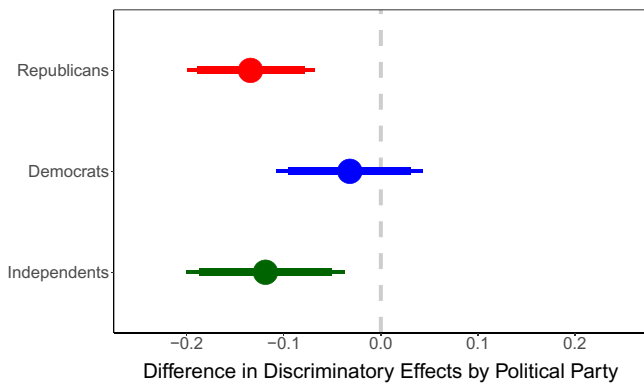
Does discrimination also manifest in how high of a quality response members of the public give? We show the results from our measure of response quality (survey completion) in *SI Appendix, Fig. S16*. Among the public, Black requesters are 9.5% less likely to receive a completed survey than all else equal White requesters. (For politicians, this number is 7.5%.) This suggests that Black people receive lower response rates and lower-quality responses.

Do the discriminatory effects vary depending on the race or ethnicity of the recipient? Fig. 2 shows our treatment effects organized by whether the recipient is White, Black, Asian, Hispanic, or some other identity category. As can be seen, the only racial group that does not discriminate against African Americans is Black people themselves. All other groups discriminate at comparable levels; White citizens and citizens of other racial minorities (when pooled together for the sake of statistical power) are just as likely to discriminate against Black people. This finding is important in that it provides strong evidence that other racial/ethnic minorities (when pooled together) also discriminate against Black people.

Do the discriminatory effects vary depending on the political affiliation of the recipient? Among the public, discrimination occurs among Republicans and Independents. Fig. 3 shows this visually. Republicans are 12% less likely ( $P = 0.00007$ ) relative to the base rate to respond to a Black sender (all else equal), and Independents are 13% less likely ( $P = 0.004$ ) relative to the base rate to respond to a Black sender (all else equal).



**Fig. 2.** Discrimination effects among the US public broken by recipient race. The effect of an African American sender (vs. a White sender) on response rates of the public. Points are coefficients (sized by  $N$ ), and bars are the respective 90% (thicker) and 95% (thinner) CIs. Effects listed here include individual fixed effects. Effects are scaled relative to the mean response rate in the control group (i.e., they are in percentage of the base rate units). Public:  $N(ALL) = 500,000$ ;  $N(W) = 352,632$ ;  $N(B) = 55,320$ ;  $N(A) = 14,344$ ;  $N(H) = 38,650$ ; and  $N(HAO) = 92,048$ . The public discriminates against Black people. This is driven by Whites and other non-Black respondents; Black people do not discriminate against Black senders.

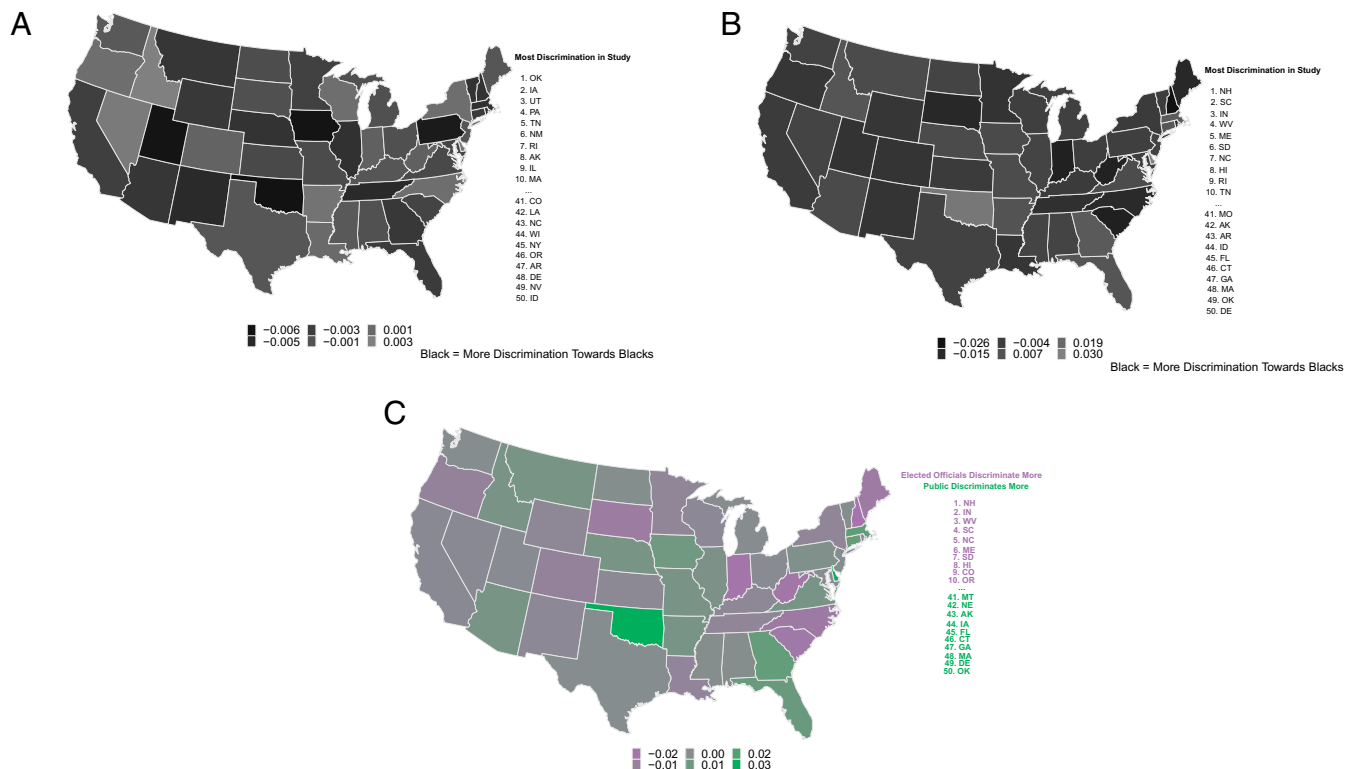


**Fig. 3.** Discrimination among the public is driven by Republicans and Independents. The figure displays the effect of an African American sender (vs. a White sender) on response rates of the public among various partisan/racial subgroups. Points are coefficients (sized by  $N$ ), and bars are the respective 90% (thicker) and 95% (thinner) CIs. Effects are scaled relative to the mean response rate in the control group (i.e., they are in percentage of the base rate units). Discrimination among the public is driven primarily by Republicans and Independents. Discrimination is larger for White Democrats than Black Democrats (*SI Appendix, Fig. S27*).

In contrast, Democrats are only 3% less likely to respond to a Black sender; that said, this effect, although negative, is not statistically significant ( $P = 0.41$ ). (The Democratic effect is significantly different from both the Republicans and Independents effects.) We have also crossed our effects by political party and race together. *SI Appendix, Fig. S28* shows this visually. There, we test differences in levels of discrimination across White and

Black people in these political parties. As can be seen, White Democrats discriminate 18.8% more (relative to the base rate) than Black Democrats. The same cannot be said of Republicans and Independents, where all races in these groups show signs of discriminating against Black email senders.

Does our measure of paper cut discrimination vary by social context? Below, we provide evidence that the discrimination we document is not concentrated in any individual region; rather, it is widespread across the United States. In this way, it is systemic; regardless of where they live, Black people are likely to face everyday discrimination. There are several ways to see this result visually. Fig. 4 plots our discrimination estimates for all states in the contiguous United States (Hawaii and Alaska are omitted for ease in visualization). Fig. 4 *A* and *B* shows effects broken down by state for the public and elected officials separately. States with darker shading see higher levels of discrimination against Black people. Fig. 4 shows that, although there is some variation in treatment effects across the country, discrimination occurs in most states. In states like Tennessee and New Mexico (where Black residents are, all else equal, 26.2% less likely to get a response than Whites), Pennsylvania (where Blacks are, all else equal, 35% less likely to get a response than Whites), Utah and Iowa (where Blacks are, all else equal, 38.7% less likely to get a response than Whites), and Oklahoma (where Blacks are, all else equal, 39.9% less likely to get a response than Whites), discrimination appears to be especially high. That said, discrimination also appears to be the norm rather than the exception. *SI Appendix, Fig. S31* shows the effect estimates shown in the map along with their corresponding 95% CIs. *SI Appendix, Fig. S31* shows that the states with the highest and lowest levels of discrimination against Black people are statistically distinct from one another, but many states in



**Fig. 4.** Paper cut discrimination by the public and elected officials across the United States. (A) Paper cut discrimination among the public. (B) Paper cut discrimination among politicians. (C) Difference between politicians and the public. Discrimination effect estimates by state. For the maps in *A* and *B*, darker colors correspond to higher levels of discrimination against Black people. In the map in *C*, purple states are where elected officials discriminate more than the public, and in green states, the public discriminates more than elected officials. Paper cut discrimination is systemic, and the modal state exhibits no difference between elected officials' and citizens' behaviors (*SI Appendix, Fig. S18*).

the middle categories are not. *SI Appendix, Fig. S24* shows this information in a slightly different way—omitting (iteratively) one state at a time. That figure shows that no individual state disrupts the general pattern of systemic discrimination among the public. As best we can tell, there are a few outlier states where discrimination against Black people is especially high; however, in a dominant majority of states, discrimination is the same.

Some may be surprised by these patterns. Indeed, many may expect that discrimination occurs mostly in the South and less so in other regions. However, this is not uniformly what other measures of racial animus/discrimination show (69), nor what we find (*SI Appendix, Figs. S38 and S39*). Our results suggest that paper cut discrimination against Black people in the United States is systemic; it is a problem that manifests itself across the country, and (in many instances) these levels of racial discrimination are more homogeneous than a simplistic South/non-South characterization would suggest.

When it comes to differences between elected officials and the public, it is possible that in some locations, elected officials are a moderating force when compared with their constituents, while in other areas, elites are more extreme than the citizens they represent. To explore this possibility, we estimate our effects among the public, elected officials, and the difference between the two (across all 50 states in the United States). Fig. 4C plots the difference between the public and elected officials (purple states are where elected officials discriminate more than the public; in green states, the public discriminates more than elected officials). As can be seen, although there are some differences in places where the public discriminates more (and vice versa), the modal result is for states to have no difference between the public and their elected officials (*SI Appendix, Figs. S18 and S29*).

Finally, in *SI Appendix*, we collected 60 observable state-level social, demographic, and political characteristics and used them to predict levels of discrimination among the public, levels of discrimination among elected officials, and the differences between the two. *SI Appendix, Figs. S14 and S15* show the correlates among the groups individually. While most of the correlations are small, interestingly one of the strongest predictors of the public's racial response patterns is the political climate of the state; discrimination against Black people is lower in states controlled by Democratic governors. This does not hold true among elected officials; Democratic- and Republican-controlled states see similar levels of racial discrimination among politicians. Moreover, *SI Appendix, Fig. S13* shows that none of the 60 characteristics we collected predict the gap in discriminatory behavior between the two groups. Any differences between politicians' and the public's behavior are outliers that do not appear to be explained by readily available empirical explanations. This provides further support to the idea that, as a general rule, the public and their elected officials tend to be aligned when it comes to racial discrimination.

## Conclusion

In this paper, we reported results from a large-scale audit study of the public. Our experimental results show that Americans discriminate against African Americans in small, seemingly negligible but vitally important everyday online interactions. Taken literally, our results show that Black people are less likely to receive an email response (and a high-quality one at that) from their fellow citizens when the email includes a request to fill out a survey. In a broader sense, however, our results reveal that Whites and members of other racial/ethnic groups (not including Black people themselves) are less inclined to volunteer their time to assist Black individuals in simple requests for help than they are to respond to such requests from people who are exactly the same in every way apart from the fact that they are White. Our

results are present across individuals of all races/ethnicities except African Americans, among all geographic regions, and within individuals of various political parties/backgrounds. Our work is important in that it explores elite- and mass-level behaviors, and the discrimination that the public exhibits shows an unsettling degree of congruence with (or if anything, heightened levels of discrimination compared with) what elected officials and some other social elites themselves do.

Our experimentally derived measure takes an important next step in understanding the nature, scope, and coverage of everyday racial discrimination in the United States. We capture a type of discrimination not previously measured. To drive home this point, we provide one additional comparison. *SI Appendix, Figs. S38 and S39* relate our measure to those used in the literature. As can be seen in these figures, our measure is largely independent from previously used metrics. The form of everyday or paper cut discrimination that we document here has not been fully documented in previous research.

Measuring racial discrimination is an area of social science that has been well trod but is, simultaneously, in need of continued rigorous research designs. Our paper offers some methodological advances and overcomes some methodological issues that have proved to be tricky in the past. Our measure has some key desirable features in that it provides 1) a behavioral measure—rather than an attitudinal one—of racial discrimination that is assessed in 2) a real-world setting—rather than in an artificial laboratory—with 3) an experimental research design that holds other factors constant among 4) a nationally representative and large pool of participants.

Our paper provides important insights into discrimination against African Americans in the United States. Particularly, our research design allows us to study an important and underappreciated form of discrimination—what we call paper cut, “everyday,” or “day-to-day” discrimination. This form of discrimination potentially occurs in simple everyday interactions—like responding to a request for information/help—that our study explicitly measures. Since the forms of social interactions that breed paper cut discrimination are often more common than those that breed more extreme discriminatory acts, documenting the nature and scope of this important form of bias provides us with a fuller understanding of the broad spectrum of racial discrimination that African Americans face.

Our work adds to the study of racial discrimination in important ways. That said, it does have some limitations. Our paper does not document all of the many variants of racial discrimination that Black people face across contexts and interactions (although we argue that no single paper can be reasonably expected to do that). Moreover, as our primary goal was the first-order task of documenting the extent, nature, and scope of paper cut discrimination, we did not tease apart all of the many potential mechanisms driving our effects. Future work would do well to dig deeper into the mechanisms driving paper cut discrimination among the public, be they motivated by taste-based or statistical considerations, whether they stem from pro-White favoritism or anti-Black animosity, whether they are due to features correlated with race, or if we can attribute them to something else entirely. Similarly, future research should examine the extent to which paper cut discrimination occurs against other racial/ethnic groups and along other socially relevant dimensions, like gender, social class, sexuality, and so on. Doing so will help tease apart potential mechanisms and give a sense of the breadth and scope of discrimination against other marginalized groups. Finally, our paper uses a specific type of helping behavior—volunteering to respond to an email invitation to take a survey—that approximates the types of interactions people have every day. However, in this paper, we have not addressed whether our effects vary by the difficulty of the request. If paper cut discrimination increases as the cost of the request for help increases, we would expect our effects to be a conservative estimate of just how much discrimination African

Americans face in the United States. Future research should test this theory.

Even with these limitations, our paper offers an important contribution in the current moment, given the heightened attention to racial injustices in the United States. Our work provides further evidence that, in addition to examining discrimination against racial minorities in the employment and public service domains, scholars should also focus on the many interpersonal interactions common to everyday life.

**Data Availability.** The data and code for this article are posted on the Harvard Dataverse at <https://doi.org/10.7910/DVNCJ7YRF> (70).

1. W. E. B. Du Bois, *The Souls of Black Folk: Essays and Reflections* (A. C. McClurg & Co, 1903).
2. W. E. B. Du Bois, *To the nations of the world* (remarks by Chairman Alan Greenspan at the Annual Dinner and Francis Boyer Lecture of The American Enterprise Institute for Public Policy Research, Washington, DC). <https://www.federalreserve.gov/boarddocs/speeches/1996/19961205.htm>. Accessed 20 June 2013.
3. S. Camp, *Closer to Freedom* (University of North Carolina Press, 2005).
4. W. E. B. Du Bois, *Black Reconstruction in America* (Routledge, 2017).
5. A. Acharya, M. Blackwell, M. Sen, *Deep Roots* (Princeton University Press, 2020), vol. 6.
6. S. Mustakeem, *Slavery at Sea* (University of Illinois Press, 2016).
7. M. Alexander, *The New Jim Crow* (The New Press, 2020).
8. J. Johnson, *The Autobiography of an Ex-Colored Man* (Penguin, 1990).
9. J. H. Franklin, E. B. Higginbotham, *From Slavery to Freedom: A History of African Americans* (McGraw-Hill, ed. 9, 2010).
10. L. Lawrence, *The Politics of Force* (University of California Press, 2000).
11. D. Knox, W. Lowe, J. Mummolo, Administrative records mask racially biased policing. *Am. Polit. Sci. Rev.* **114**, 619–637 (2020).
12. A. White, Misdemeanor disenfranchisement? The demobilizing effects of brief jail spells on potential voters. *Am. Polit. Sci. Rev.* **113**, 311–324 (2019).
13. D. Neumark, R. Bank, K. Van Nort, Sex discrimination in restaurant hiring. *Q. J. Econ.* **111**, 915–941 (1996).
14. M. Bertrand, S. Mullainathan, Are Emily and Greg more employable than Lakisha and Jamal? *Am. Econ. Rev.* **94**, 991–1013 (2004).
15. L. Quillian, D. Pager, O. Hoxel, A. H. Midtbøen, Meta-analysis of field experiments shows no change in racial discrimination in hiring over time. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 10870–10875 (2017).
16. D. Butler, D. Broockman, Do politicians racially discriminate against constituents? *Am. J. Pol. Sci.* **55**, 463–477 (2011).
17. D. Butler, *Representing the Advantaged* (Cambridge University Press, 2014).
18. M. Costa, How responsive are political elites? *J. Exp. Polit. Sci.* **4**, 241–254 (2017).
19. A. Harris, M. Sen, Bias and judging. *Annu. Rev. Polit. Sci.* **22**, 241–259 (2019).
20. R. Block, What about disillusionment? Exploring the pathways to black nationalism. *Polit. Behav.* **33**, 27–51 (2011).
21. R. Block Jr., Backing Barack because he's black: Racially motivated voting in the 2008 election. *Soc. Sci. Q.* **92**, 423–446 (2011).
22. I. K. White, When race matters and when it doesn't: Racial group differences in response to racial cues. *Am. Polit. Sci. Rev.* **101**, 339–354 (2007).
23. N. A. Valentino, V. L. Hutchings, I. K. White, Cues that matter: How political ads prime racial attitudes during campaigns. *Am. Polit. Sci. Rev.* **96**, 75–90 (2002).
24. C. Davenport, *Media Bias, Perspective, and State Repression: The Black Panther Party* (Cambridge University Press, 2009).
25. V. L. Hutchings, H. Walton Jr., A. Benjamin, The impact of explicit racial cues on gender differences in support for confederate symbols and partisanship. *J. Polit.* **72**, 1175–1188 (2010).
26. V. L. Hutchings, C. Wong, Racism, group position, and attitudes about immigration among blacks and whites 1. *Du Bois Rev.* **11**, 419–442 (2014).
27. V. L. Hutchings, A. E. Jardina, Experiments on racial priming in political campaigns. *Annu. Rev. Polit. Sci.* **12**, 397–402 (2009).
28. V. L. Hutchings, N. A. Valentino, The centrality of race in American politics. *Annu. Rev. Polit. Sci.* **7**, 383–408 (2004).
29. A. J. Berinsky, V. L. Hutchings, T. Mendelberg, L. Shaker, N. A. Valentino, Sex and race: Are black candidates more likely to be disadvantaged by sex scandals? *Polit. Behav.* **33**, 179–202 (2011).
30. E. Washington, Do majority-black districts limit blacks' representation? The case of the 1990 redistricting. *J. Law Econ.* **55**, 251–274 (2012).
31. E. O. Ananat, E. Washington, Segregation and black political efficacy. *J. Public Econ.* **93**, 807–822 (2009).
32. E. Washington, How black candidates affect voter turnout. *Q. J. Econ.* **121**, 973–998 (2006).
33. P. Essed, *Understanding Everyday Racism: An Interdisciplinary Theory* (Sage, 1991), vol. 2.
34. C. Pierce, Stress analogs of racism and sexism. *Ment. Health (Lond.)* **33**, 277–293 (1995).
35. D. W. Sue et al., Racial microaggressions in everyday life: Implications for clinical practice. *Am. Psychol.* **62**, 271–286 (2007).
36. D. Pager, H. Shepherd, The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets. *Annu. Rev. Sociol.* **34**, 181–209 (2008).
37. J. Guryan, K. K. Charles, Taste-based or statistical discrimination: The economics of discrimination returns to its roots. *Econ. J.* **123**, F417–F432 (2013).
38. F. E. Aboud, The formation of in-group favoritism and out-group prejudice in young children: Are they distinct attitudes? *Dev. Psychol.* **39**, 48–60 (2003).
39. K. Marquis, S. Marquis, M. Polich, Response bias and reliability in sensitive topic surveys. *J. Am. Stat. Assoc.* **81**, 381–389 (1986).
40. S. Pfaff, C. Crabtree, H. L. Kern, J. B. Holbein, Do street-level bureaucrats discriminate based on religion? A large-scale correspondence experiment among American public school principals. *Public Adm. Rev.* **81**, 244–259 (2021).
41. N. Carnes, J. Holbein, Do public officials exhibit social class biases when they handle casework? Evidence from multiple correspondence experiments. *PLoS One* **14**, e0214244 (2019).
42. A. White, N. Nathan, J. Faller, What do I need to vote? *Am. Polit. Sci. Rev.* **109**, 129–142 (2015).
43. H. J. G. Hassell, J. B. Holbein, M. R. Miles, There is no liberal media bias in which news stories political journalists choose to cover. *Sci. Adv.* **6**, eaay9344 (2020).
44. D. M. Butler, C. Crabtree, Moving beyond measurement: Adapting audit studies to test bias-reducing interventions. *J. Exp. Polit. Sci.* **4**, 57–67 (2017).
45. D. Butler, C. Crabtree, "Audit studies in political science" in *Advances in Experimental Political Science*, J. Druckman, D. Green, Eds. (Cambridge University Press, Cambridge, United Kingdom, 2021), pp. 42–55.
46. C. Crabtree, K. Dhima, Auditing ethics: A cost-benefit framework for audit studies. *Polit. Stud. Rev.* **10.1177/14789299211046153** (2021).
47. C. Crabtree, *An Introduction to Conducting Email Audit Studies in Audit Studies: Behind the Scenes with Theory, Method, and Nuance* (Springer, 2018), pp. 103–117.
48. M. Gaddis, How black are Lakisha and Jamal? Racial perceptions from names used in correspondence audit studies. *Sociol. Sci.* **4**, 469–489 (2017).
49. C. Crabtree, "Measuring and Explaining Discrimination," PhD thesis, University of Michigan, Ann Arbor, MI (2019).
50. E. Hersh, *Hacking the Electorate* (Cambridge University Press, 2015).
51. B. Fraga, *The Turnout Gap* (Cambridge University Press, 2018).
52. J. Holbein, D. S. Hillygus, *Making Young Voters* (Cambridge University Press, 2020).
53. C. Kennedy et al., "Comparing survey sampling strategies" *Pew Research Center* (2021). <https://www.pewresearch.org/methods/2018/10/09/comparing-survey-sampling-strategies-random-digit-dial-vs-voter-files/>. Accessed 8 December 2021.
54. D. Butler, C. Volden, A. Dynes, B. Shor, Ideology, learning, and policy diffusion: Experimental evidence. *Am. J. Pol. Sci.* **61**, 37–49 (2017).
55. A. Gelman, J. Carlin, Beyond power calculations: Assessing type s (sign) and type m (magnitude) errors. *Perspect. Psychol. Sci.* **9**, 641–651 (2014).
56. K. Imai et al., Estimating treatment effect heterogeneity in randomized program evaluation. *Ann. Appl. Stat.* **7**, 443–470 (2013).
57. A. Gelman, The connection between varying treatment effects and the crisis of unreplicable research: A Bayesian perspective. *J. Manag.* **41**, 632–643 (2015).
58. A. Coppock, T. J. Leeper, K. J. Mullinix, Generalizability of heterogeneous treatment effect estimates across samples. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 12441–12446 (2018).
59. G. Charness, U. Gneezy, M. A. Kuhn, Experimental methods: Between-subject and within-subject design. *J. Econ. Behav. Organ.* **81**, 1–8 (2012).
60. A. Coppock, Avoiding post-treatment bias in audit experiments. *J. Exp. Polit. Sci.* **6**, 1–4 (2019).
61. J. Montgomery, B. Nyhan, M. Torres, How conditioning on posttreatment variables can ruin your experiment and what to do about it. *Am. J. Pol. Sci.* **62**, 760–775 (2018).
62. J. Esarey, N. Danneberg, A quantitative method for substantive robustness assessment. *Political Sci. Res. Methods* **3**, 95–111 (2015).
63. J. L. Kalla, D. E. Broockman, Campaign contributions facilitate access to congressional officials: A randomized field experiment. *Am. J. Pol. Sci.* **60**, 545–558 (2016).
64. K. Baicker et al.; Oregon Health Study Group, The Oregon experiment—Effects of Medicaid on clinical outcomes. *N. Engl. J. Med.* **368**, 1713–1722 (2013).
65. J. Haushofer, J. Shapiro, The short-term impact of unconditional cash transfers to the poor: Experimental evidence from Kenya. *Q. J. Econ.* **131**, 1973–2042 (2016).
66. J. Holbein, Childhood skill development and adult political participation. *Am. Polit. Sci. Rev.* **111**, 572–583 (2017).
67. R. Chetty, N. Hendren, L. F. Katz, The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *Am. Econ. Rev.* **106**, 855–902 (2016).
68. J. Holbein, Discrimination against Black and Hispanic Americans is highest in hiring and housing contexts: A meta-analysis of correspondence audits. <https://doi.org/10.7910/OSF.IO/4EQN2>. Deposited 1 December 2021.
69. C. W. Smith, R. J. Kreitzer, F. Suo, The dynamics of racial resentment across the 50 US states. *Perspect. Polit.* **18**, 527–538 (2020).
70. R. Block Jr., C. Crabtree, J. B. Holbein, J. Q. Monson, Replication Data for: Are Americans less likely to reply to emails from Black people relative to White people? Harvard Dataverse. <https://doi.org/10.7910/DVNCJ7YRF>. Deposited 1 December 2021.