



Federal Reserve Bank of Chicago

**Identifying Taste-Based Discrimination:
Effect of Black Electoral Victories on
Racial Prejudice and Economic Gaps**

Jung Sakong

May 17, 2021

WP 2021-07

<https://doi.org/10.21033/wp-2021-07>

**Working papers are not edited, and all opinions and errors are the responsibility of the author(s). The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.*

Identifying Taste-based Discrimination: Effect of Black Electoral Victories on Racial Prejudice and Economic Gaps

Jung Sakong*

May 17, 2021

Abstract

I test for the causal impact of Black electoral victories in local elections on White Americans' attitude toward Black Americans. Using Race Implicit Attitude Test scores as a measure of racial prejudice and close-election regression-discontinuity design for causal inference, I find Black electoral victories cause measures of racial bias to rise, by 4% of the average Black-White difference in IAT scores. Simultaneously, they widen racial gaps in unemployment and mortgage denials. Interpreting these close electoral victories as instrumental variables, I find a large causal effect of prejudice-based racial discrimination on Black-White economic gaps.

*Federal Reserve Bank of Chicago (email: jung.sakong@chi.frb.org). These views are those of the authors and do not reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System. I am extremely grateful to my committee chair Amir Sufi and members Marianne Bertrand, Raghuram Rajan and Luigi Zingales for their continuous guidance and support. I also thank Jane Dokko, Lancelot Henry de Frahan, Lars Hansen, Paymon Khorrami, Helen Koshy, Canice Prendergast, Dan Sullivan as well as seminar participants at the University of Chicago Applied Microeconomics Lunch, the Stigler Lunch Workshop and the Federal Reserve System Applied Microeconomics Conference. All errors are my own.

1 Introduction

The election in 2008 and two-term presidency of Barack Obama motivated discussions as to how having a Black leader affects White Americans' attitude toward Black Americans in general.¹ Empirical analysis of Black leaders has been challenging due to two issues: First, relative to population, Black politicians are disproportionately few in numbers, and second, Black politicians are not elected randomly. This paper seeks to overcome those issues by looking at a much broader set of local elections of Black politicians, and – aided by the larger sample size – by using close elections to establish causality from electing Black leaders to racial attitudes and actual racial gaps in economic outcomes.

Black candidates' electoral victories can affect racial attitudes of White people through multiple mechanisms, with different effects. On one hand, the election of Black politicians could dispel negative stereotypes associated with Black Americans.² On the other hand, the election of successful Black politicians may lead to a backlash effect among White Americans who perceive such success as a threat to their dominance.³ Furthermore, having elected a Black politician, White Americans may feel more justified in holding and expressing racial bias via the self-licensing effect (Monin and Miller (2001)).⁴

In this paper, I estimate a causal impact of Black politicians' electoral victories on White Americans' racial prejudice and Black Americans' relative economic outcomes, using local elections and a close election regression discontinuity design.

¹See Parker (2016) for a comprehensive review and DellaVigna (2010) for an economic analysis. Empirical studies find results in both directions, depending on sample and methodology (for example, Goldman and Mutz (2014); Kinder and Chudy (2016); Tesler and Sears (2010); Tesler (2016)).

²Beaman et al. (2009); Beaman et al. (2012) and Chattopadhyay and Duflo (2004) find that randomized rise of female politicians reduce gender bias and lead to positive outcomes for girls in India.

³Theories of intergroup bias based on social dominance argue that Black politicians may raise the perceived threat to White Americans' racial dominance and motivate them to engage in backlash (Olzak (1990); Sidanius et al. (1992)). While traditional sociological theories of the threat hypothesis typically studied changes in Black population shares, more recent psychological studies find higher intergroup apathy when the outgroup competes with their high status (Harris et al. (2008); Cikara (2015)). In contrast to knowledge-based stereotypes, prejudice derived from threat and modulated by guilt tends to be more non-cognitive (March et al. (2018)).

⁴Guilt in response to intergroup wrongs has long been understood to lead to regulating one's own prejudice (Allport (1954); Amodio et al. (2007)). Monin and Miller (2001) argue that one act of racial tolerance endows the individual with a moral license to express biased views in other dimensions, exhibiting the so-called self-licensing effect. Effron et al. (2009) find that endorsing Barack Obama during the 2008 presidential election made his supporters more likely to display racial bias against Black Americans. Kouchaki (2011) further finds that even knowing of others' un-prejudiced behavior leads individuals to freely express biased views.

In order to measure racial prejudice, I use Race Implicit Attitude Test (IAT) score data from Project Implicit Database, compiled by Xu et al. (2014). The IAT is a widely used test of subconscious or implicit racial bias (Greenwald et al. (1998)). Project Implicit Database collected voluntary online IAT surveys, along with demographic and other respondent information. To focus on actual changes in racial attitudes rather than changes in the composition of who took the test, I project out demographic fixed effects from the raw IAT scores, assign weights based on sex, education, and age groups to be representative, and finally use only responses by test-takers who took the IAT to fulfill an assignment from school or work (the “mandatory” sub-sample). I then aggregate White respondents’ IAT scores to a county-month panel. Given about 250,000 completed surveys per year, the resulting panel data have informative local variation and correlate with other proxies for racial prejudice and Black-White economic gaps.

Data on local elections are scraped from Our Campaigns, a collaborative website that compiles electoral information.⁵ I use data on any US election with a sub-state constituency. This includes US federal House representatives, mayors, city council members, county executives, and county council members. I extend the set of politicians beyond representatives and mayors in order to increase the sample size because the representation of Black politicians in these races is so low. For candidates whose race is not identified in Our Campaigns, I use surnames and facial recognition to indicate race, admittedly both imperfect methods.

To overcome the identification challenge that electoral outcomes are not random, I use a standard close-election regression-discontinuity (RD) design. Among areas that had a close election where the top two candidates include one Black and one White candidate, I compare the average IAT scores between areas where a Black candidate won against areas where a Black candidate lost, in the three-year period after the elections. The three-year period before an election is used for placebo tests. I define a “close election” as one with a margin of less than 10% vote between the winner and the runner-up, consistent with the close-election literature.⁶ (Lee and Lemieux (2010)). The results are robust to alternative definitions of close elections, in particular to defining them with narrower margins as well as to data-driven definitions (for example to optimal bandwidths proposed by Imbens and Kalyanaraman (2012) and Calonico et al. (2014)), which imply

⁵More information available online, <https://www.ourcampaigns.com/about.html>.

⁶For example, Vogl (2014) uses a 15% bandwidth in examining close interracial mayoral elections.

wider margins. Results are also robust to other variations in the RD specifications.

I find that a Black politician's victory causes racial prejudice to rise as measured by the IAT score. The average effect size is around 0.02, which is about 4-5% of the raw average gap in IAT scores between White and Black Americans, where higher IAT score indicates more racial bias. Before the election, there is no difference in IAT scores in areas with a Black winner versus a White winner. The rise in prejudice is more pronounced in areas with higher Black population share and lower income, consistent with the perceived-threat mechanism.

Simultaneously, using the same RD design, Black politicians' electoral victories cause Black workers to transition more into non-employment and causes Black applicants to be denied mortgages more frequently – consequently getting lower mortgage amounts originated per capita – relative to White workers and applicants. Within the employment results, the effect is concentrated among firms in the 0-19 employee-count category. Most of these small firms are not subject to anti-discrimination enforcement by the Equal Employment Opportunity Commission (EEOC); that the effect is concentrated among those firms suggests that the EEOC enforcement may have a positive effect in preventing a change in racial bias from causing a tangible gap in employment outcomes.

In the last part of the paper, I interpret the close election of a Black politician as a time-varying instrumental variable (IV) for local racial prejudice. This IV strategy is a novel approach to identifying “taste-based discrimination” in the data (Becker (1971)), that is, discriminatory behavior driven by preferences toward one group (e.g. White people) over another group (e.g. Black people). I highlight how this approach of identifying discrimination differs from what has been done in the literature so far. To make sense of the various data components, it is helpful to think of close electoral victories of Black candidates as the forcing variable or instrument (i.e., Z), racial prejudice measures as data on the independent variable of interest or the outcome of the first-stage (i.e., X), and interracial economic disparity measures as data on the ultimate or second-stage outcome variable (i.e., Y). The estimated second-stage impact of racial bias on racial economic gaps is large. For example, a naive back-of-the-envelope calculation for mortgage denials using the second-stage estimate suggests that the majority of the Black-White gap may be due to racial discrimination.

This paper makes three broad contributions. First, I document the causal direction of what has been dubbed the “Obama effect” by showing that close electoral victories lead to higher racial

prejudice among White people against Black people. Second, I propose using such a time-varying “shock” to racial prejudice to identify prejudice-driven – or taste-based – racial discrimination in employment and mortgage markets. Third, I construct demographically adjusted local panels of IAT scores as a local measure of racial prejudice – validated against other racial gaps – and a local election data set with the candidate races identified, which are novel data sets that have broader applications and are publicly available. Studying Black politicians has been difficult due to small samples, especially with close elections, and this new data set can be useful in future research.

Using close electoral victories of Black candidates, the IV methodology can be applied to study the impact of racial prejudice or taste-based discrimination in a variety of different settings, such as in education, policing, and judicial systems. In this paper, I focus on labor and mortgage market discrimination, where we have the most robust economics research to lean on as well as anti-discrimination policies already in place. Nonetheless, even in these markets, I find that the increases in prejudice widen economic gaps. The methodological benefits of “shocking” prejudice go beyond just having another instrument to identify the effect of race. This approach also gets at racial prejudice or taste-based discrimination directly, as opposed to overall racial disparity that can be driven by other mechanisms such as statistical discrimination.

2 Racial Prejudice: Implicit Association Test (IAT)

Since Becker (1971), economists have taken as given an unobserved “taste for discrimination” and derived implications for observed behaviors, which are then tested with data. A key innovation of Charles and Guryan (2008) was to bring in empirical proxies for the “taste for discrimination” directly, using survey responses from the General Social Survey (GSS). This paper builds on this idea by measuring the “taste for discrimination” with the IAT, which measures implicit (i.e., sub-conscious) racial prejudice using an online response module. This section describes in detail the IAT, the Project Implicit Database that measured IAT scores across the US (Xu et al. (2014)), and constructing local measures of IAT and their time variation.

The aim of the data construction is to be able to measure changes in the racial attitudes of a local area, in the absence of data on repeated test scores at the individual level. The idea that the underlying racial attitudes can change over time is not new. A large psychological literature

studies interventions that can reduce both implicit and explicit racial bias in lab experiments (Lai et al. (2014); Lai et al. (2016)). Becker (1971) himself wrote, “Another proximate determinant is geographical and chronological location: discrimination may vary from country to country, from region to region within a country, from rural to urban areas within a region, and from one time period to another.”

2.1 IAT Data from Project Implicit

The Race IAT is a widely used test of racial bias (Greenwald et al. (1998)). Respondents are shown either pictures of faces (Black or White) and words (associated with good or bad). They use the same set of buttons on the keyboard to classify the faces into Black or White categories and words into good or bad categories. IAT is based on the premise that if a respondent has a stronger association in their mind between being White and being good, the classification exercise will take longer when they have to use the same button to classify a face as being Black and a word as being good. The main measure, the D score, is the difference in time it takes to classify when Black faces and good words are paired together (i.e., use the same button) versus when Black faces and bad words are paired together. Higher scores indicate more racial bias.

Since its inception in Greenwald et al. (1998), there has been an explosion of studies using the IAT to test various implicit attitudes. Recent meta-studies find that the Race IAT is a good predictor of racial discrimination (Oswald et al. (2013)). The IAT has also been used in economics as well as other social sciences to measure racial bias (Reuben et al. (2014); Lowes et al. (2015); Lowes et al. (2017); Carlana (2019); Chetty et al. (2020)). Chetty et al. (2020), in particular, use the same data source from the Project Implicit Database.

I get data on online IAT scores from the Project Implicit Database (Xu et al. (2014)). Figure 1b shows the distribution of the raw IAT scores in the dotted line. The data span years 2003-2017, with roughly 250,000 tests completed per year. The online survey also collects demographic (age, sex, race, education), political, and religious information; explicit racial bias questions comparable to those asked in the General Social Survey (GSS); and self-reported zip code.

The rich demographic and other details self-reported in the Project Implicit Database show how racial prejudice as measured by the IAT score co-varies with demographics (see Table 1). By far the most salient variable is race: Black respondents’ average IAT score is -0.04 compared

with 0.39 for non-Hispanic White respondents, that is, non-Hispanic White respondents are more likely to have negative implicit attitudes toward Black people. To a much smaller magnitude, IAT scores are higher for men and those with less education, two widely cited co-variates for racial bias (Hodson and Busseri (2012)). Income and wealth gradients also show up, but the magnitude is small, as is the fraction of respondents reporting their income and wealth. Political ideology also correlates highly with the IAT scores. Lastly, implicit bias correlates strongly with explicitly stated racial preferences, here measured as thermology (i.e., “do you feel warmer toward...”) and directly elicited preference for Black people or White people (Table 1).

Since the online surveys are voluntary, the sample is self-selected and highly unrepresentative of the US population.⁷ Such self-selection is a feature of most measures of amorphous entities such as racial prejudice. For example, Google searches for racial slurs is another measure of racial bias that is affected by who chooses to use Google and for that purpose (Stephens-Davidowitz (2014)). The self-selection issue is potentially larger in the case of the IAT because it is a test explicitly designed for racial bias.

Self-selection of test-takers introduces several issues to be addressed. First, the sample is not representative of the broader population, nor especially about the geographical areas they represent – a well known issue among psychologists who use the database. Second and more importantly, my empirical methodology compares outcomes between areas with a close Black electoral winner and those with a close loser. If different test-takers self-select in response to the close election, that can induce a bias in the estimate. For example, if following a close electoral victory of a Black candidate, more biased individuals take the test for the first time, I would observe an increase in the area’s average level of bias, but it would be entirely due to the composition of test-takers, rather than an increase in the level of bias.⁸

To address the first issue, I re-weight observations to be more representative of the corresponding area’s distribution of sex, age, and education according to the American Community Survey, as is often done in the psychology literature. I further project out dummies for sex, age, and education groups over the entire sample. The presence of observable demographics also helps address the second issue. As shown in Table 1, bias levels correlate with some of these observable char-

⁷Online Appendix Table A.1a

⁸I discuss each of these issues in more detail in the Online Appendix.

acteristics. If the composition of test-takers changes in response to the election result, I expect that to show up in these observable groups. Therefore, by checking that controlling for the demographic characteristics does not change the estimate, I can verify that composition likely did not change (Altonji et al. (2005)). Last and most importantly, I compare responses of those who were assigned to take the test against those who took it voluntarily. How responses differ between the more and less self-selected subgroups within the sample will be informative in evaluating how the self-selection of the sample as a whole may bias my results. Using the question, “What brought you to this website?” I classify a response as being mandatory if the respondent chose “Assignment for work” or “Assignment for school.” The average IAT score among White respondents is slightly higher for mandatory test-takers.⁹

2.2 IAT Local Panel Construction and Validation

I construct county-month or county-year local panels of raw or adjusted average IAT scores.¹⁰ For the raw averages, I simply take all IAT scores among White respondents and average them for each county-month or county-year.¹¹

In the “adjusted” IAT measure, I first project White respondents’ IAT scores on age, race, education (nine buckets), gender, and experimental variables (month, hour, weekday, and order of experiment), and take the residuals. I then aggregate these residuals to the county-month or county-year level, weighing each observation according to the fully interacted gender, age, and education shares in the 2008-2012 American Community Survey (ACS).¹²

Figure 1a plots the time series average of these adjusted responses for White respondents, with shaded standard errors and number of responses in bars (right axis). Against an overall downward trend, there is a visible rebound starting around 2009, overlapping with the first presidential term of Barack Obama.

In this section, I further validate the local-area measures of racial prejudice. Aggregating to

⁹For more details on the respondents’ reasons for taking the test and their average characteristics and IAT scores, see Online Appendix Table A.1b.

¹⁰While respondents report their zip codes – which are finer than counties – the zip codes are not available in the public version of the data for privacy reasons.

¹¹While it is interesting to think about the racial attitudes of non-White respondents, the sample is smaller and it is outside the scope of this paper.

¹²See the Online Appendix and Online Appendix Table A.1a for more details on the weights.

a county panel reduces a lot of the variation as can be seen in Figure 1b. The benefit, however, is that the IAT averages can then be studied in conjunction with events (e.g., elections) and other outcomes (e.g., economic gaps). The bottom panel of Table 1 shows the adjusted R^2 from regressing the individual-level IAT scores on geography $\times$ time-period dummies. For the adjusted scores, the county-month variation takes up 28% of the individual-level variation, which is sizable for geographic factors. Figure 1c plots county-level average adjusted scores, using the standard break-points (e.g., 0.65-1.10 for high implicit bias against Black people).

I further validate the information content of the local IAT panel in three ways. First, in a cross-sectional county-level correlation, areas with higher IAT averages also have Black workers with lower earnings than White workers, higher transitions to unemployment during the time period, and getting more mortgage rejections but taking out larger mortgage loans per capita (see the top two panels of Table 2a). Here, the relationships are only correlations, but these economic gaps will be used as second-stage outcomes later in the paper. Second, areas with higher level of IAT averages also have more searches for anti-Black racial slurs and “KKK” on Google (Stephens-Davidowitz (2014)), higher likelihood of corporal punishment at school for Black students than White students in states where spanking is legal, and a higher level of survey-elicited prejudice in the General Social Survey (GSS) following Charles and Guryan (2008). These are other measures of racial prejudice used in the literature, which correlate with the areas’ average IAT.¹³

Lastly, Acharya et al. (2016) proposed historical slavery intensity in 1860 as an instrument for today’s level of racial prejudice, further instrumenting for slavery-intensity using cotton production. Table 2b uses their empirical methodology and shows that areas with more intensive slavery in 1860 have higher IAT averages today among White respondents.¹⁴

3 Local Election Data

Political elections provide the ideal setting for studying the effect of a salient exemplar for both substantive and methodological reasons. Substantively, politicians are highly visible, partly because they have to campaign to attract votes. Methodologically, election outcomes can be un-

¹³See the Online Appendix for more details on these prejudice measures.

¹⁴These are county-level regressions comparing within-state. See the Online Appendix for more details. The slavery-IAT correlation is also documented in Payne et al. (2019). For the relationship between racial prejudice and the legacy of slavery, see also Nunn (2009).

certain and provide an identification strategy to study the causal effect of having an exemplar.

3.1 Data from Our Campaigns

Data on local electoral outcomes come from Our Campaigns, a collaborative (Wikipedia-like) website that aggregates electoral information.¹⁵ I use data on any US election with a sub-state constituency. This includes US federal House representatives, mayors, city council members, county executives, and county council members. Given the disproportionately low representation of Black politicians, extending the set of politicians beyond the commonly considered representatives and mayors drastically increases the sample size. Since 2003, there were 202 Black congressional victories in the US House of Representatives, 12 Black gubernatorial victories, and 215 Black mayoral victories in the Our Campaigns data using the racial identification scheme described below (Table 3a). Focusing on close elections between Black and White candidates, defining close elections as those with a vote margin of 10% or less between the top two candidates, leaves 68 elections and 39 victories for Black candidates among the three commonly studied, most visible offices. Including state legislatures, county councils, and city councils, along with other elected municipal offices (e.g., county president) increases these numbers six- to seven-fold, to 501 close Black-White elections and 247 victories for Black candidates in such elections. The sample-size gain is traded off against the lower visibility and salience of these local offices in determining statistical power.

To compile the election data used, I scraped the Our Campaigns website for the following information: most recent map of the jurisdiction and the associated longitude-latitude coordinates,¹⁶ history of electoral races with date, type, candidates and their vote counts and shares, and candidate information, in particular the user-supplied tags for race and photos for facial identification of candidates' race.¹⁷ Since I am expanding the set of politicians studied beyond what is commonly done in the literature, identifying candidates' races is a challenge, which I address below.

¹⁵www.ourcampaigns.com

¹⁶Electoral jurisdictions are mapped to counties using shapefiles, and overlapping proportions are calculated using areas. Details are in the Online Appendix.

¹⁷See example website interface in Online Appendix Figures A.5a (for map), A.5a (for history) and A.5c (for candidate information).

3.2 Identifying Candidates' Race

Candidates' race information is provided for a small subset of those on the website, so I use two other methods to identify candidates' race. Whenever the website contains direct information on race, I use that information. This small subset is also useful for judging the accuracy of the other two methods. Table 3b shows how the candidates classified according to tags are classified using the other two methods of surnames and facial recognition. Note that there may be errors in the user-supplied tags too.

The second source of racial identification is the candidates' last name.¹⁸ Surnames are a widely used source of demographic information, not only the US.¹⁹ Using the Census 2000 and 2010 surnames files (Word et al. (2008); Comenetz (2016)), I classify a candidate as belonging to one racial group, if 80% or more of people with that candidates' last name belong to that group. For example, candidates with the surname Little are classified as more likely to be Black, while candidates with surname Hansen are classified as more likely to be White. Surnames are highly informative for Hispanic and Asian Americans, but less so for Black and White Americans who tend to share common surnames. This can be seen in the lower success rate of classifying Black and White candidates to their races in Table 3b.²⁰

The last source of racial identification is the candidates' photo. Whenever a candidate has an associated photo, I use facial recognition to infer the candidate's race (see Fu et al. (2014) for summary of techniques). To this end, I use free software provided by Face++. Their algorithm classifies race into Asian, Black, and White. I again check the accuracy of facial recognition by comparing the race from face recognition with the tag information. Among the candidates classified as White according to the facial-recognition algorithm, 81% are tagged as being White on Our Campaigns; among those classified as Black according to the algorithm, 88% are tagged as being Black (Table 3b). Given the noise, I use the other sources of information first if they are available, before resorting to facial recognition. Nonetheless, facial recognition yields many instances where race can be identified especially for Black politicians for whom surnames are less

¹⁸First names can also be informative about race, but there is no comprehensive tabulation of first names' racial association as there exists for surnames.

¹⁹For example, Neggers (2018) used surnames in India to identify religious and caste identity.

²⁰Audit studies have exploited strong racial associations of first names (Bertrand and Mullainathan (2004)), but no comprehensive list of first names and racial associations exists.

informative and Face++ is able to make more accurate classifications, as can be seen in Table 3b. Even with the noise in classification, facial recognition is useful for the sample gain in this study and for creating a database of politicians by race more broadly.

When racial information from these three sources disagree, I rely first on the tag information, then surname, then facial recognition.

4 Empirical Methodology and Impact on Prejudice

To introduce the empirical methodology, I begin with the following “ideal experiment.” The goal is to isolate a “shock” to White residents of a county, who often make up the majority of its population and more often are disproportionately in positions of influencing economic decisions. In the ideal experiment, the county gets an associated office-holder whom the residents identify as Black. I want to then measure the racial prejudice of those White residents solely in response to having a Black political leader.

There are two main challenges to this ideal experiment. First, Black leaders are not elected randomly. I adopt the close-election regression-discontinuity design widely used in the literature (Lee and Lemieux (2010); Eggers et al. (2015)). Second, the elected official may adopt policies that directly affect economics outcomes. I will address the second concern in Section 5 where I use the close electoral victory as an instrumental variable.

An alternative research design is a difference-in-difference estimator, comparing the outcomes before and after election in the counties influenced, compared with those around them that are outside the jurisdiction, assuming visibility of the elected leader is lower outside their jurisdiction. The main concern with this approach is negative selection: Black politicians are more likely to run and to win in areas where majority White voters’ racial bias against Blacks is decreasing (Marshall and Ruhil (2006)). This will bias my estimate to find that areas where a Black politician comes into office will experience a decrease in racial prejudice.²¹

Indeed, a difference-in-difference estimate finds zero impact of Black electoral victory on racial bias (Online Appendix Figure A.7a and Table A.2). I compare counties affected by a Black electoral victory with surrounding counties in the same state, for three years after the election versus

²¹ An alternative solution is to add parametric time controls for trends in racial prejudice (e.g., linear trends). In this case, identification relies having the correct parametric specification.

three years preceding the election. I follow Gormley and Matsa (2011) in constructing a difference-in-difference estimator with multiple events.²² To overcome the endogeneity concern, I turn to a regression-discontinuity design, exploiting close elections.

4.1 Regression Discontinuity Design and Estimates

Since Lee et al. (2004) and Lee (2008), a large literature has utilized close-election regression-discontinuity (RD) design as an identification scheme to estimate the causal impact of a political victory, with well-established econometrics methodology.²³ According to this literature, winning a closely contested election is considered as good as random (Eggers et al. (2015)).

I use various RD specifications and find consistent estimates. In the main specification, I follow the close-election literature and define “close elections” as those where the vote margin between the winner and the runner-up is 10% or less. The choice of bandwidth is of first-order importance in a valid RD design. My choice of a 10% bandwidth is slightly narrower than what has been used in the literature.²⁴ It is also narrower than the optimal robust bandwidth (Calonico et al. (2014)), and I report results using that wider bandwidth as well as other specifications. Lastly, using the 10% definition, roughly half of Black-White close elections are won by the Black candidate in each type of election (see last two columns of Table 3a), consistent with the unpredictability of close elections.

For each election, I compile counties fully or partially in the electoral jurisdiction for three years following the election for the main estimate. The final sample compiles samples across all elections where the top two candidates included one Black and one White candidate. “Vote margin” is defined as the vote share of the Black candidate minus that of the White candidate, i.e., positive “vote margin” indicates a Black electoral victory. For observations with election i , county j and event month t , I estimate the following regression model:

$$Y_{ijt} = \alpha + \gamma_1 1\{\text{vote margin} > 0\}_{it} + \delta_0 [\text{vote margin}]_{it}^- + \delta_1 [\text{vote margin}]_{it}^+ + \eta_{ijt}$$

²²The Online Appendix contains details of the difference-in-difference estimate, for both IAT and for economic gaps.

²³Imbens and Lemieux (2008); Lee and Lemieux (2010); Imbens and Kalyanaraman (2012); Calonico et al. (2014)

²⁴For example, Vogl (2014) used 15% bandwidth in examining interracial mayoral elections.

where the first indicator is a dummy for Black electoral victories, and the second and third terms are linear controls for negative and positive vote margins, respectively. Observations are weighted only by the share of county population impacted by the election.

Because the identifying assumption is that omitted variables cannot predict who will win among closely contested elections (Lee and Lemieux (2010)), the RD estimator is consistent whether I use just the level of IAT in the after period or the change in IAT from before to after the election. I use the level here, and use the three years preceding the election as a placebo test. That is, if electoral victories in close elections are random, election victory should not predict previous prejudice levels.

Columns (1)-(2) and (7)-(8) of Table 4 show the regression estimates. The top panel is the main estimate from the after-election period, and the bottom panel is the placebo estimate from the before period. Columns (1) and (7) use the raw, unadjusted IAT averages as the outcome variable, for all White respondents and all White mandatory test-takers, respectively. Columns (2) and (8) use the adjusted IAT averages, for all and mandatory test-takers. Column (2) is the main estimate, with a statistically significant effect size of around 0.02. This specification is also graphically shown in Figure 2: At the discontinuity threshold at 0, there is no gap in the linear regression lines before the election, but it shows up after the election.

Relative to close electoral defeats, close electoral victories of Black candidates causally raise White residents' IAT scores by around 0.02 on average in a county. The effect size is large, at 4-5% of the gap in IAT scores between all Black and White respondents in the Project Implicit Database. For comparison, in psychological lab experiments, one of the more effective interventions to reduce implicit bias – implementation intentions – reduces an experimental subject's IAT by 0.05 over three days.²⁵ Given that this is an estimate of an individually targeted intervention over a three-day period, an increase in IAT scores of around 0.02 in county averages over multiple quarters is a large causal impact on implicit bias.

By contrast, the corresponding pre-period regressions yield estimates that are statistically not different from zero. By using the same outcome in the before period, I confirm that Black electoral victories among close elections are indeed “shocks.” Moreover, I am able to run the same RD specification quarter-by-quarter relative to the election quarter, and I plot the RD estimates in

²⁵Lai et al. (2013); Lai et al. (2014); Lai et al. (2016); Paluck et al. (2020)

Figures 3a and 3b using raw and adjusted IAT scores, respectively. There is a stable zero difference leading up to the election, but discontinuity rises in the quarters following the election, and comes back down eventually. Both the pre-period placebo and the time series plots of RD estimates suggest that the RD design is picking up a causal estimate.

The RD estimate is robust to alternative specifications of the RD design. Since bandwidth is often the most sensitive parameter of an RD design (Imbens and Lemieux (2008); Imbens and Kalyanaraman (2012)), Column (3) of Table 4 uses a narrower bandwidth at 5%. Columns (4)-(6) and (9) follow robust bandwidths of Calonico et al. (2014), which are wider than the 10% in the main specification and as wide as 32%, along with triangular kernels, i.e., elections with larger margins are assigned lower weights. Columns (5) and (6) use quadratic and cubic polynomials, respectively, on either side of the discontinuity threshold. Across these specifications, the post-period estimate is consistent at around 0.02, with statistically insignificant estimates before the election.

4.2 Heterogeneity by County and Office Types

The treatment effect is concentrated almost exclusively in counties with a higher-than-median population share of Black residents and counties with lower-than-median income level, and to a lesser extent in counties with a higher-than-median initial level of prejudice²⁶ (first three columns in Table 5). These estimates are from a regression specification that additionally interacts the indicator for an electoral victory of a Black candidate and the linear terms on either side of the discontinuity threshold with a dummy for the county having a higher level of IAT, population share of Black residents, or income, independently.²⁷

The county-type-dependent results are broadly consistent with the threat mechanism (Olzak (1990)), which postulates that racial bias among White people is a non-cognitive response to a perceived threat represented by an out-group, due to competing numbers or status. Consistent with this mechanism, the bias response to Black candidates' victories is higher in areas with more Black residents. Also consistent, the response is higher in lower-income counties, where weaker

²⁶Initial level is measured as the average adjusted IAT level in 2005, the first year with more than 100,000 White responses.

²⁷RD design is already a taxing empirical design in terms of statistical power, and these heterogeneity-by-county-type estimates effectively cut the number of counties used to estimate each parameter.

economic security among White residents may lead them to perceive a higher threat from gains in status of Black residents via electoral victories (Bobo (1988)).

The bias response is also higher in areas with higher initial levels of prejudice, albeit with weaker statistical power, and the level of prejudice is positively correlated with Black population share and negatively with income (Table 1). This may help explain the apparent contrast with Beaman et al. (2009), which found women politicians have a positive impact on women's outcomes in India in more open-minded provinces. Possibly more importantly, perceived threat – the likely mechanism at play in this paper – is proposed as one of the main explanations for racial bias but less so for gender bias, for which cognitive stereotypes may be the dominant mechanism as pointed out by Beaman et al. (2009).

The last three columns of Table 5 present sub-sample estimates by the type of office for mayor and US House of Representative elections. These offices are more salient and visible, and the estimated impact on IAT is higher, at 0.06 and 0.12, respectively. However, the sample sizes are small, and the standard errors are large, especially for US House elections.

5 Effect on Racial Gaps in Employment and Credit

With close electoral victories of Black politicians increasing White residents' racial bias, this section finds a simultaneous impact on economic gaps between Black people and White people. There are many possible outcomes, but I focus on employment and mortgage markets for two reasons. First, these markets and racial discrimination therein are the most heavily researched with widely used data sets. Second, these markets already have policies in place to prevent racial discrimination. If I find rising taste-based discrimination here, the impact may be greater in areas without such policy protection.

Using the same RD design, I find a large impact on widening economic outcomes. This real impact contrasts with psychology lab experiments that can lower implicit bias but often do not find corresponding behavioral responses, possibly because it is difficult for lab studies to follow up with substantial real-life behaviors (Lai et al. (2016)). By contrast, in studies that rely on non-experimental data such as this paper, it is more difficult to establish clean causation onto racial prejudice, but it is easier to measure behavioral impact; this contrast highlights the benefit of being

able to identify a shock to racial prejudice in non-experimental settings.

5.1 Data on Employment and Mortgages

For employment and mortgage outcomes, both data sources and variable definitions are widely used in the literature. The main outcome of interest is defined as outcomes for Black workers and loan applicants minus outcomes for White workers and loan applicants.²⁸ Data on economic outcomes are in county-year panels, to be consistent across the two data sources.²⁹

Data on employment come from Quarterly Workforce Indicators (QWI) and refer to all private employment at the county-level (Abowd et al. (2009)). The QWI county panel data are aggregated from the Longitudinal Employer-Household Dynamics linked employer-employee microdata. The LEHD data cover 95% of private-sector jobs, and are constructed from state-level administrative records (such as Unemployment Insurance earnings data) and federal administrative data. I use two main measures. The first is the share of those employed in the beginning of the year who transition into non-employment during that year. This is a flow measure and is likely the most responsive to underlying attitudes, given that wages may be sticky in the labor market. The second measure is the employment-to-population ratio, which is the corresponding stock measure.

Data on mortgage denial come from the Home Mortgage Disclosure Act (HMDA) data (HMDA (2017)). I use two main measures. The first is the Black-White difference in mortgage denial rates, calculated as the number of mortgage applications for owner-occupied housing that were rejected, divided by all loans that were originated or denied. The rejection rate is a widely used metric in the mortgage discrimination literature, but rejection rate is affected by who applies – an endogenous decision – so I also look at the Black-White difference in log mortgage origination amounts per capita.

5.2 Regression Discontinuity Estimates

To assess the impact of Black electoral victories on racial gaps in economic outcomes, I use the same RD specification on the election-county-year data. Again, relying on close elections achieves

²⁸This is equivalent to regressing the outcomes themselves on the RD regressors interacted with the corresponding race indicators.

²⁹The highest observable frequency of the public HMDA data is year.

identification in an environment where non-close electoral victories of Black candidates can be confounded by improving racial attitudes or Black Americans' economic condition. Furthermore, since identification comes from the assumption that close victories are effectively random, additional controls (e.g., income in mortgage regressions) can improve the efficiency of the estimator but do not impact its consistency (Lee and Lemieux (2010), section on "Irrelevance of including baseline covariates").

Looking at economic gaps, the top panel of Table 6 shows that Black politicians' electoral victories cause Black workers to transition more into non-employment (flow), fewer Black workers to be in employment (stock), Black mortgage applicants to be denied mortgages at a higher rate, and Black people to get lower amounts of mortgages per capita than their White counterparts.

The magnitudes are again large. Following a close electoral victory for a Black candidate, Black workers are 2.5 percentage points more likely to transition to unemployment per year in flow terms, resulting in a 5 percentage point lower employment-to-population ratio in the three-year time period following the election on average. Simultaneously, the rejection rate for Black mortgage applicants rises 3.5 percentage points relative to White applicants, leading to a 25 percentage-point lower per capita mortgage origination in relative terms. These are large increases in racial gaps that accompany the rise in racial bias.

In the next section, I will combine these RD estimates on economic gaps along with the changes in racial bias as measured by IAT scores to quantify how much of these economic racial gaps are caused by racial bias.

5.3 Interaction with Anti-discrimination Law

There are regulations at many levels of the government that are designed to prevent discriminatory practices in employment and mortgage markets. The Equal Employment Opportunity Commission (EEOC) is at the center of this effort for employment, processing the majority of employment-related discrimination claims (Selmi (1996)). Federal enforcement of the EEOC applies to private firms with 15 or more employees.

The QWI employment data provide statistics by firm-size categories in the following groups: firms with 0-19, 20-49, 50-249, 250-499, and 500 or more employees. Computing the transition-to-unemployment ratios for firms in each firm-size category independently and taking the Black-

White difference as above, I show in the bottom panel of Table 6 the RD estimates for these employee-count categories.

The positive impact of Black electoral victories on the Black-White gap in transition to unemployment is concentrated among firms in the 0-19 employee count, for the majority of which the EEOC does not enforce. The estimate is 20% bigger than the overall estimate. However, with the conservative RD methodology, heterogeneity estimates are not precise.³⁰

These estimates are consistent with the possibility that the anti-discriminatory enforcement by the EEOC may be effective in stopping the rise in racial bias against Black workers from translating into higher firing of Black employees. Another possibility is that taste for discrimination is decreasing in firm size; this latter possibility is less likely given that the estimates – while imprecisely estimated – do not exhibit any monotonic pattern.

6 Identifying Taste-based Discrimination

Close electoral victories of Black candidates increases both racial bias as measured by the IAT and economic gaps between White people and Black people. If the only substantial impact of the electoral victory is on White residents' perceived racial bias, I can interpret the close electoral victories as instrumental variables for the racial bias, to estimate an IV estimate of racial bias on economic gaps.

This IV methodology depends crucially on the assumption that the political victories do not otherwise affect economic outcomes, especially through public policies enacted by Black politicians.

Here I rely on a large political science literature examining the impact of Black mayors and their policies. The key question in these studies is whether active political participation matters for voicing the preferences of minority constituents, as opposed to passive participation in the form of voting; in particular they test whether Black mayors enact more liberal policies.³¹ Most recent studies find that the racial identity of mayors has no direct policy impact due to constraints that

³⁰These heterogeneity estimates are nonetheless more precise than the IAT heterogeneity estimates. Here, the heterogeneity is on the outcome variable, whereas with the IAT heterogeneity by region, an additional estimate halves the number of events used.

³¹Keller (1978); Abney and Hutcheson Jr (1981); Saltzstein (1989); Marschall and Ruhil (2006); Hopkins and McCabe (2012); Nye et al. (2014); Eggers et al. (2015); Brollo and Troiano (2016)

municipal executives face, except in a few policing practices.³² These studies suggest that if the election of Black politicians has an economic impact, it is less likely to be through actual public policy changes. And that elections with Black candidates receive more attention as measured by higher White and Black voter turnout suggests a possible psychological channel if Black electoral victory affects economic outcomes of Black and White constituents (Washington (2006)).

6.1 Instrumental Variable (IV) Estimator

I start from the standard RD exclusion restriction that whether the Black candidate wins or loses in a close election is uncorrelated with local conditions, except through its effect on local racial prejudice: $E[1\{\text{vote margin} > 0\}_{it} \varepsilon_{ijt} | \text{IAT}_{ijt}] = 0$ where ε_{ijt} is the unobservable factors in $Y_{ijt} = \beta_0 + \beta_1 \text{IAT}_{ijt} + \varepsilon_{ijt}$ for some Black-White economic disparity measure Y_{ijt} in county j in year t , surrounding an election i .

This exclusion restriction motivates the following two-stage least squares (2SLS) estimator.

$$\begin{aligned} Y_{ijt} &= \beta_0 + \beta_1 \hat{\text{IAT}}_{ijt} + \tilde{\delta}_0 [\text{vote margin}]_{it}^- + \tilde{\delta}_1 [\text{vote margin}]_{it}^+ + \varepsilon_{ijt} \\ \text{IAT}_{ijt} &= \alpha + \gamma_1 1\{\text{vote margin} > 0\}_{it} + \delta_0 [\text{vote margin}]_{it}^- + \delta_1 [\text{vote margin}]_{it}^+ + \eta_{ijt} \end{aligned}$$

for election i , county j and year t , where the linear controls on either side of the vote margin enter both first stage and second stage regressions. The instrumental variable is $1\{\text{vote margin} > 0\}_{it}$, which has a first-stage partial F-stat of 8.34, making it a borderline weak instrument.

The main IV estimation results are reported in the middle panel of Table 6. Given the exclusion restriction, the first row estimates show the causal impact of 1-point increase in the composition-adjusted IAT score among White residents in the county. To interpret the magnitude of the estimate for mortgage rejection rate, a 0.1 increase in racial bias among White residents as measured by the IAT would lead the Black-White rejection rate gap to widen by 6.3 percentage points.

For a naive back-of-the-envelope calculation, the mortgage rejection rate estimate of 0.634 along with the overall average IAT score of 0.31 implies that if all racial bias against Black residents as measured by the IAT fell to 0, the rejection rate gap would fall by roughly 20 percentage points,

³²Studies looking at party or gender find similar results (Ferreira and Gyourko (2014); Ferreira and Gyourko (2009)).

equivalent to the actual rejection rate gap. Such an extrapolation likely over-estimates the true counterfactual, given endogenous responses and that the estimated β_1 is a local average treatment effect (LATE), but nevertheless demonstrates that the IV estimate is large.

6.2 Relation to Other Identification Schemes

This paper proposes using close electoral victories of Black candidates as a time-varying “shock” to racial prejudice to identify directly taste-based racial discrimination in observational data. Here, I discuss how this identification scheme relates to and differs from other identification schemes for racial discrimination.

In broad conceptual terms, the identification schemes for racial discrimination in observational data and experiments estimate one of the following regressions:

$$Y_i = \beta \text{Black}_i + X_i \Gamma + \tilde{\varepsilon}_i \quad (1a)$$

$$Y_i = \beta \hat{\text{Black}}_i + \varepsilon_i \quad (1b)$$

$$Y_{ij} = \beta \text{Black}_i \times \text{prejudice}_j + \delta \text{Black}_i + \varepsilon_{ij} \quad (1c)$$

$$Y_{ijt} = \beta \text{Black}_i \times \hat{\text{prejudice}}_{jt} + \delta \text{Black}_i + \varepsilon_{ijt} \quad (1d)$$

for individuals i , their economic outcome Y_i (e.g. whether their mortgage application is rejected) and Black_i is individual i ’s race indicator, for example. Hats indicate instrumented second-stage regressors.

The majority of studies on racial discrimination using observational data estimate regressions in the form of equation 1a (Charles and Guryan (2011)), where X_i is a vector of controls such as income. The estimated coefficient $\hat{\beta}$ identifies racial discrimination if controls remove all variations unrelated to discrimination. There is a well-known over-controlling problem here: if an area is highly prejudiced in both the mortgage market and the labor market, Black residents may both earn less than White residents and have their mortgage applications rejected more frequently. In such a setting, it is possible for the specification to estimate $\hat{\beta} = 0$, even though the mortgage market also practices racial discrimination. In other words, $\hat{\beta}$ from such specifications can only tell us how discriminatory the market for Y_i is relative to the markets for the controls in X_i .

Audit studies get around this issue by essentially “shocking” Black_i by experimentally varying the perceived race of the applicant. By submitting identical applications but only varying the name of the applicant, these studies estimate equation 1b, where $\hat{\text{Black}}_i$ can be thought of as the perceived probability that an applicant is Black given that their name is Lakisha or Jamal as opposed to Emily or Greg (Bertrand and Mullainathan (2004)).

A few papers take a different approach. Starting from the insight that the taste for discrimination is a combination of the target of discrimination and the intensity of racial prejudice, they estimate regressions of the form 1c by using local measures of prejudice $_j$. Charles and Guryan (2008) used General Social Survey responses to get geographical variations in prejudice $_j$. Other papers extend this by “shocking” the level of prejudice: Acharya et al. (2016) instrument for prejudice $_j$ use the historical prevalence of slavery driven by cotton production intensity in 1860, and Schindler and Westcott (2021) using the historical presence of American soldiers (GIs) in Europe.

This paper extends this second approach. I associate whole county-years with a level of racial prejudice and find proxies for it. I “shock” this local prejudice with close electoral victories. Then, I estimate equation 1d for a consistent estimate of $\hat{\beta}$ without additional controls. The panel structure allows me to verify that my instrument is truly a “shock” by using the previous period for placebo tests.³³

7 Discussion and Conclusion

This paper makes three contributions. First, by using IAT scores from the Project Implicit Database to measure racial bias and close-election RD design for identification, I find that the election of Black candidates raises White residents’ racial bias against Black residents. Second, using this as an instrumental variable for shifting “tastes for discrimination” and looking at Black-White economic gaps in employment and mortgage outcomes, I find a large effect of racial prejudice on Black-White gaps, implying a substantial role for racial discrimination in explaining observed racial gaps. Third, by web-scraping Our Campaigns for election data and using facial identification to categorize racial groups of political candidates, I construct detailed local-election data with

³³Conceptually this method compares different geographical areas as opposed to different markets (for example, mortgage market compared to labor market as above). As a result, I cannot conclude which market in a geographical area exhibits prejudice, but I can say that city X has more discrimination than city Y.

candidate race information, which can be useful for studying active political participation across demographic groups.

This research raises several policy implications. First, more research is needed into what can change racial prejudice and how that can impact economic outcomes. Second, the EEOC may be effective in preventing a rise in prejudice from spilling over into actual unemployment for Blacks, based on the result that the adverse effects are concentrated among small firms that the EEOC does not reach. While there are costs to be considered expanding EEOC to small employers, this finding suggests the benefit may be bigger than expected especially if racial bias is prone to change. Third, this research finds a large role of racial discrimination in employment and mortgage markets: More anti-discriminatory policy may be needed for both efficiency and equity.

Using Black electoral victories in close elections as a “shock” to racial bias can be applied to study the impact of racial prejudice in other settings, such as in education, policing, and the justice system, to name a few. More broadly, natural experiments that “shock” racial prejudice can be a useful methodology for directly identifying taste-based discrimination in observational data.

References

- Abney, F. G. and Hutcheson Jr, J. D. (1981). Race, representation, and trust: Changes in attitudes after the election of a black mayor. *Public Opinion Quarterly*, 45(1):91–101.
- Abowd, J. M., Stephens, B. E., Vilhuber, L., Andersson, F., McKinney, K. L., Roemer, M., and Woodcock, S. (2009). The lehd infrastructure files and the creation of the quarterly workforce indicators. In *Producer dynamics: New evidence from micro data*, pages 149–230. University of Chicago Press.
- Acharya, A., Blackwell, M., and Sen, M. (2016). The political legacy of american slavery. *The Journal of Politics*, 78(3):621–641.
- Allport, G. W. (1954). *The nature of prejudice*. Addison-Wesley.
- Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of political economy*, 113(1):151–184.
- Amodio, D. M., Devine, P. G., and Harmon-Jones, E. (2007). A dynamic model of guilt: Implications for motivation and self-regulation in the context of prejudice. *Psychological Science*, 18(6):524–530.
- Beaman, L., Chattopadhyay, R., Duflo, E., Pande, R., and Topalova, P. (2009). Powerful women: does exposure reduce bias? *The Quarterly journal of economics*, 124(4):1497–1540.
- Beaman, L., Duflo, E., Pande, R., and Topalova, P. (2012). Female leadership raises aspirations and educational attainment for girls: A policy experiment in india. *science*, 335(6068):582–586.
- Becker, G. S. ([1957] 1971). *The Economics of Discrimination*. University of Chicago press.
- Bertrand, M. and Mullainathan, S. (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American economic review*, 94(4):991–1013.
- Bobo, L. (1988). Group conflict, prejudice, and the paradox of contemporary racial attitudes. In *Eliminating racism*, pages 85–114. Springer.
- Brollo, F. and Troiano, U. (2016). What happens when a woman wins an election? evidence from close races in brazil. *Journal of Development Economics*, 122:28–45.

- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.
- Carlana, M. (2019). Implicit stereotypes: Evidence from teachers' gender bias. *The Quarterly Journal of Economics*, 134(3):1163–1224.
- Charles, K. K. and Guryan, J. (2008). Prejudice and wages: an empirical assessment of becker's the economics of discrimination. *Journal of Political Economy*, 116(5):773–809.
- Charles, K. K. and Guryan, J. (2011). Studying discrimination: Fundamental challenges and recent progress. *Annu. Rev. Econ.*, 3(1):479–511.
- Chattopadhyay, R. and Duflo, E. (2004). Women as policy makers: Evidence from a randomized policy experiment in india. *Econometrica*, 72(5):1409–1443.
- Chetty, R., Hendren, N., Jones, M. R., and Porter, S. R. (2020). Race and economic opportunity in the united states: An intergenerational perspective. *The Quarterly Journal of Economics*, 135(2):711–783.
- Cikara, M. (2015). Intergroup schadenfreude: Motivating participation in collective violence. *Current opinion in behavioral sciences*, 3:12–17.
- Comenetz, J. (2016). Frequently occurring surnames in the 2010 census.
- DellaVigna, S. (2010). The obama effect on economic outcomes: Evidence from event studies. Technical report, Citeseer.
- Effron, D. A., Cameron, J. S., and Monin, B. (2009). Endorsing obama licenses favoring whites. *Journal of experimental social psychology*, 45(3):590–593.
- Eggers, A. C., Fowler, A., Hainmueller, J., Hall, A. B., and Snyder, J. M. (2015). On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races. *American Journal of Political Science*, 59(1):259–274.
- Ferreira, F. and Gyourko, J. (2009). Do political parties matter? evidence from us cities. *The Quarterly journal of economics*, 124(1):399–422.
- Ferreira, F. and Gyourko, J. (2014). Does gender matter for political leadership? the case of us mayors. *Journal of Public Economics*, 112:24–39.
- Fu, S., He, H., and Hou, Z.-G. (2014). Learning race from face: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 36(12):2483–2509.
- Goldman, S. K. and Mutz, D. C. (2014). *The Obama effect: How the 2008 campaign changed white racial attitudes*. Russell Sage Foundation.
- Gormley, T. A. and Matsa, D. A. (2011). Growing out of trouble? corporate responses to liability risk. *The Review of Financial Studies*, 24(8):2781–2821.
- Greenwald, A. G., McGhee, D. E., and Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*, 74(6):1464.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., and Banaji, M. R. (2009). Understanding and using the implicit association test: Iii. meta-analysis of predictive validity. *Journal of personality and social psychology*, 97(1):17.
- Harris, L. T., Cikara, M., and Fiske, S. T. (2008). Envy as predicted by the stereotype content model: Volatile ambivalence. *Envy: Theory and research*, pages 133–147.
- HMDA (2017). Download historic hmda data.
- Hodson, G. and Busseri, M. A. (2012). Bright minds and dark attitudes: Lower cognitive ability predicts greater prejudice through right-wing ideology and low intergroup contact. *Psychological science*, 23(2):187–195.
- Hopkins, D. J. and McCabe, K. T. (2012). After it's too late: estimating the policy impacts of black mayoralities in us cities. *American Politics Research*, 40(4):665–700.
- Imbens, G. and Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *The Review of economic studies*, 79(3):933–959.
- Imbens, G. W. and Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of econometrics*, 142(2):615–635.
- Keller, E. J. (1978). The impact of black mayors on urban policy. *The Annals of the American Academy of Political and Social Science*, 439(1):40–52.
- Kinder, D. and Chudy, J. (2016). After obama. In *The Forum*, volume 14, pages 3–15. De Gruyter.
- Kouchaki, M. (2011). Vicarious moral licensing: the influence of others' past moral actions on moral behavior. *Journal of personality and social psychology*, 101(4):702.

- Kurdi, B., Seitchik, A. E., Axt, J., Carroll, T., Karapetyan, A., Kaushik, N., Tomezsko, D., Greenwald, A. G., and Banaji, M. R. (2018). Relationship between the implicit association test and intergroup behavior: A meta-analysis.
- Lai, C. K., Hoffman, K. M., and Nosek, B. A. (2013). Reducing implicit prejudice. *Social and Personality Psychology Compass*, 7(5):315–330.
- Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J.-E. L., Joy-Gaba, J. A., Ho, A. K., Teachman, B. A., Wojcik, S. P., Koleva, S. P., et al. (2014). Reducing implicit racial preferences: I. a comparative investigation of 17 interventions. *Journal of Experimental Psychology: General*, 143(4):1765.
- Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., Calanchini, J., Xiao, Y. J., Pedram, C., Marshburn, C. K., et al. (2016). Reducing implicit racial preferences: Ii. intervention effectiveness across time. *Journal of Experimental Psychology: General*, 145(8):1001.
- Lee, D. S. (2008). Randomized experiments from non-random selection in us house elections. *Journal of Econometrics*, 142(2):675–697.
- Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of economic literature*, 48(2):281–355.
- Lee, D. S., Moretti, E., and Butler, M. J. (2004). Do voters affect or elect policies? evidence from the us house. *The Quarterly Journal of Economics*, 119(3):807–859.
- Lowes, S., Nunn, N., Robinson, J. A., and Weigel, J. (2015). Understanding ethnic identity in africa: Evidence from the implicit association test (iat). *American Economic Review*, 105(5):340–45.
- Lowes, S., Nunn, N., Robinson, J. A., and Weigel, J. L. (2017). The evolution of culture and institutions: Evidence from the kuba kingdom. *Econometrica*, 85(4):1065–1091.
- March, D. S., Gaertner, L., and Olson, M. A. (2018). On the prioritized processing of threat in a dual implicit process model of evaluation. *Psychological Inquiry*, 29(1):1–13.
- Marschall, M. J. and Ruhil, A. V. (2006). The pomp of power: Black mayoralties in urban america. *Social Science Quarterly*, 87(4):828–850.
- Monin, B. and Miller, D. T. (2001). Moral credentials and the expression of prejudice. *Journal of personality and social psychology*, 81(1):33.
- Neggers, Y. (2018). Enfranchising your own? experimental evidence on bureaucrat diversity and election bias in india. *American Economic Review*, 108(6):1288–1321.
- Nunn, N. (2009). The importance of history for economic development. *Annu. Rev. Econ.*, 1(1):65–92.
- Nye, J. V., Rainer, I., and Stratmann, T. (2014). Do black mayors improve black relative to white employment outcomes? evidence from large us cities. *The Journal of Law, Economics, & Organization*, 31(2):383–430.
- Olzak, S. (1990). The political context of competition: Lynching and urban racial violence, 1882–1914. *Social Forces*, 69(2):395–421.
- Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., and Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of iat criterion studies. *Journal of personality and social psychology*, 105(2):171.
- Paluck, E. L., Porat, R., Clark, C. S., and Green, D. P. (2020). Prejudice reduction: Progress and challenges. *Annual Review of Psychology*, 72.
- Parker, C. S. (2016). Race and politics in the age of obama. *Annual Review of Sociology*, 42:217–230.
- Payne, B. K., Vuletich, H. A., and Brown-Iannuzzi, J. L. (2019). Historical roots of implicit bias in slavery. *Proceedings of the National Academy of Sciences*, 116(24):11693–11698.
- Reuben, E., Sapienza, P., and Zingales, L. (2014). How stereotypes impair women’s careers in science. *Proceedings of the National Academy of Sciences*, 111(12):4403–4408.
- Saltzstein, G. H. (1989). Black mayors and police policies. *The journal of Politics*, 51(3):525–544.
- Schindler, D. and Westcott, M. (2021). Shocking racial attitudes: black gis in europe. *The Review of Economic Studies*, 88(1):489–520.
- Selmi, M. (1996). The value of the eeoc: Reexamining the agency’s role in employment discrimination law. *Ohio St. LJ*, 57:1.
- Sidanius, J., Devereux, E., and Pratto, F. (1992). A comparison of symbolic racism theory and social dominance theory as explanations for racial policy attitudes. *The Journal of Social Psychology*, 132(3):377–395.

- Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, 118:26–40.
- Tesler, M. (2016). *Post-racial or most-racial?: Race and politics in the Obama Era*. University of Chicago Press.
- Tesler, M. and Sears, D. O. (2010). *Obama's race: The 2008 election and the dream of a post-racial America*. University of Chicago Press.
- Vogl, T. S. (2014). Race and the politics of close elections. *Journal of Public Economics*, 109:101–113.
- Washington, E. (2006). How black candidates affect voter turnout. *The Quarterly Journal of Economics*, 121(3):973–998.
- Word, D. L., Coleman, C. D., Nunziata, R., and Kominski, R. (2008). Demographic aspects of surnames from census 2000. *Unpublished manuscript*, Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download>.
- Xu, K., Nosek, B., and Greenwald, A. (2014). Psychology data from the race implicit association test on the project implicit demo website. *Journal of Open Psychology Data*, 2(1).

Figure 1: IAT

Panel (a) plots the average Race Implicit Association Test (IAT) D score for White respondents from Project Implicit Database. The average has been taken after projecting out dummies for age (each whole number), nine education buckets and gender. Panel (b) plots the distribution of raw Race IAT D score from Project Implicit Database for 2002-2017. The D score has a possible range of -2 to +2, where higher number indicates bias against Black Americans. The vertical red lines indicate break points for a common description of pro-White bias.





Figure 2: Regression Discontinuity

Both panels plot bin scatters of average Race IAT D score among White respondents from Project Implicit Database: panel (a) is for the three-year window before an election (i.e. placebo period); panel (b) is for the three-year window after the election (i.e. treatment period). Each point represents within-bin averages of Race IAT D scores across election-county-month level observations), where the bins are defined by each whole-number percent of vote margin. The sample has been restricted to elections where the winner and the runner-up include one Black and one White candidate; vote margin is defined as the percentage point difference between the Black candidate and the White candidate (i.e. positive vote margins represent elections where the Black candidate won, while negative vote margins represent elections where the Black candidate lost to the White candidate). The averages were formed using Race IAT D scores after projecting out dummies for age (each whole number), nine education buckets and gender. The vertical dotted line divides Black candidates' losses on the left to their victories on the right. The thick dark-blue lines are local polynomial lines on either side of the discontinuity, along with the associated 95% confidence intervals in thin Black lines, weighted using an Epanechnikov kernel and 2% bandwidth. The vertical distance along the vertical dotted line between the two red lines is a rough estimate of the regression discontinuity estimate. For alternative specifications, see Online Appendix Figures A.8, A.9 and A.10.

(a) Before the election ("Placebo")



(b) After the election ("Treatment")





Figure 3: Time series of discontinuity

Both panels plot time series of regression discontinuity estimates of the Race IAT D score between areas with a close Black winner and a close Black loser, for non-overlapping 3-month windows relative to the election. Each red point is the estimate γ^τ from

$$Y_{ij\tau} = \alpha^\tau + \gamma^\tau 1\{\text{vote margin} > 0\}_{i\tau} + \delta_0^\tau [\text{vote margin}]_{i\tau}^- + \delta_1^\tau [\text{vote margin}]_{i\tau}^+ + \varepsilon_{ij\tau}$$

for election i , county j and relative quarter τ relative to the election quarter, with 95% confidence intervals shaded in gray. The vertical red line at 0 indicates the 3-month period starting from the month of election; points to the right of the vertical line represents non-overlapping 3-month periods after the election. The regression discontinuity estimates have been estimated from by using elections where the vote margin between the winner and the runner-up was within 10%, with linear controls on either side of the discontinuity. Regressions have been run on observations at the election-county-month level, with each observation weighted by the fraction of the county affected by the election.

(a) Raw IAT



(b) Adjusted IAT




Table 1: Project Implicit Database Summary Statistics

This table reports summary statistics for IAT D scores from the Project Implicit Database. The top panel reports the raw mean and standard deviation of the IAT D score by sub-groups, along with the share of total sample represented by that sub-group in the Project Implicit Database. The shares do not add up to 1 where there are missing/unknown values for the sub-group variables. The bottom panel reports the adjusted R^2 values regressing the IAT D variable on geographic and/or time fixed effects. "Adjusted" IAT D scores are residuals from regressing the raw IAT D scores on sex, age and education dummies, which are then weighed to be representative of the US according to weights from the American Community Survey.

Group	Mean	S.D.	Share	Group	Mean	S.D.	Share
Total	0.31	0.44		<19 (Whites)	0.38	0.40	0.22
Asian	0.33	0.43	0.05	19-22	0.41	0.40	0.25
Black	-0.04	0.45	0.11	23-29	0.38	0.41	0.19
Hispanic	0.28	0.43	0.10	30-39	0.37	0.43	0.15
White	0.39	0.41	0.65	40+ years old	0.38	0.44	0.18
Whites aged 23+ only below							
Voluntary test	0.35	0.44	0.31	Male	0.40	0.43	0.46
Mandatory	0.38	0.42	0.28	Female	0.36	0.42	0.54
Northeast	0.38	0.43	0.18	Up to HS grad	0.42	0.42	0.04
Midwest	0.38	0.42	0.26	Some college	0.41	0.41	0.28
South	0.39	0.42	0.32	College grad	0.38	0.43	0.29
West	0.35	0.43	0.24				
Income <\$20k	0.31	0.43	0.01	Wealth <\$5k	0.30	0.44	0.02
\$20k - \$40k	0.32	0.44	0.02	\$5k - \$20k	0.32	0.43	0.02
\$40k - \$60k	0.32	0.44	0.02	\$20k - \$100k	0.33	0.44	0.02
\$60k - \$90k	0.32	0.45	0.02	\$100k - \$500k	0.33	0.45	0.02
>\$90k	0.33	0.45	0.03	>\$500k	0.34	0.46	0.01
Conservative	0.44	0.41	0.26	Christian	0.40	0.42	0.45
Neutral	0.40	0.41	0.19	Other rel	0.35	0.44	0.09
Liberal	0.33	0.43	0.52	Not religious	0.35	0.43	0.33
By explicit bias measures (Whites aged 23+)							
Warm to Black	0.20	0.46	0.07	Prefer Black	0.13	0.46	0.03
Neutral	0.34	0.42	0.60	Equally	0.33	0.42	0.52
Warm to White	0.49	0.40	0.30	Prefer White	0.47	0.41	0.35
Geographic factors adjusted R^2 (Whites aged 23+)							
IAT variable	State $\times$ year		County	County $\times$ yr		County $\times$ mo	
Raw	0.01		0.01	0.01		0.01	
Adjusted	0.02		0.05	0.15		0.28	

Table 2: IAT validation

These tables report the cross-sectional and panel relationship between the local average of the Race IAT D score and the variables in column headers. The cross-section is either a county or a Nielsen Designated Market Area. See text for definition and source for each variable used here. In panel (b), the sample is restricted to slave states in 1860, and state fixed effects are included, following Acharya et al. (2016). Standard errors are in parentheses.

(a) Cross-sectional correlations

Outcome	IAT D	(s.e.)	Constant	(s.e.)	Adj R^2	N	Unit
Labor market							
Earning	-0.054	(0.018)	-0.401	(0.010)	0.008	3,064	county
To non-emp	0.015	(0.004)	-0.025	(0.001)	0.068	3,072	county
Mortgage							
Reject rate	0.052	(0.008)	0.171	(0.003)	0.063	2,511	county
Origination	0.063	(0.054)	-1.139	(0.024)	0.271	2,758	county
Other prejudice measures							
Google n-word	0.551	(0.210)	3.190	(0.032)	0.178	204	DMA
Google KKK	0.299	(0.191)	2.508	(0.045)	0.224	207	DMA
Spanking (b-w)	0.013	(0.006)	0.004	(0.001)	0.009	105	DMA
GSS	0.319	(0.266)	-0.884	(0.028)	0.061	350	county
Miscellaneous							
Obama	-0.127	(0.024)	0.504	(0.009)	0.068	3,086	county
Income	-0.146	(0.041)	3.344	(0.017)	0.068	3,098	county
Pop Black	0.103	(0.014)	0.120	(0.007)	0.007	3,103	county

(b) By historical slavery intensity in 1860

	IAT D							
	OLS		OLS		IV		1st	
Slave share	0.020	(0.010)	0.029	(0.011)	0.041	(0.014)		
Pop Black			-0.020	(0.013)	-0.026	(0.014)	0.415	(0.071)
Cotton prod							0.169	(0.027)
Constant	0.092	(0.003)	0.096	(0.003)			0.145	(0.011)
Adjusted R^2	0.039		0.039		-0.002		0.602	
N	1114		1111		1111		1115	

Table 3: Our Campaigns Summary Statistics

Panel (a) shows the number of elections falling into each category, for years 2003 and onward. Row titles indicate which type of office the elections in that row are for. The first two columns are for all electoral races; the third and fourth columns are for electoral races in which the winner and the runner-up include one White and one Black candidate; the last two columns are such electoral races between Black and White candidates that are close. Close election is defined here as those with a vote margin between the winner and the runner-up of 10% or less of votes. Odd columns show counts of all such elections; even columns show the count of elections in which the winner is a Black candidate. The last row with the total counts show that for close elections between a Black and a White candidate, the probability that the Black candidate wins is roughly half. Panel (b) shows racial categorization of candidates using the three methods.

(a) Election counts by type (Black candidate victory in parentheses)

Office type		All elections		Black-White		Close B-W	
Federal	House	2,933	(202)	111	(78)	11	(4)
State	Governor	308	(12)	26	(11)	15	(9)
	Senate	9,362	(655)	471	(156)	51	(24)
	House	33,046	(2,426)	1,518	(486)	145	(78)
	State (Other)	1,428	(92)	147	(49)	88	(37)
Local	County (Other)	7,350	(475)	448	(146)	77	(34)
	City (Other)	4,309	(601)	277	(148)	72	(35)
	Mayor	2,755	(215)	166	(74)	42	(26)
Total		61,491	(4,678)	3,164	(1,148)	501	(247)

(b) Candidate race identification

Method	Asian	Black	Hispanic	White	No info
OurCampaigns tag (count)	39	276	187	921	49,938
Surname (count)	358	151	1,824	27,160	21,854
Frac same as tag	1.00	0.86	0.95	0.91	
Facial recognition(count)	1,435	2,148		24,960	22,804
Frac same as tag	0.30	0.88		0.81	
Final candidate count	957	1,876	1,875	36,958	9,681

Table 4: Regression discontinuity: IAT D scores

Both panels display the main results from the regression discontinuity design. The main specification in columns (1) and (2) imposes a 10% bandwidth (i.e. regression sample only includes elections for which the vote margin is at most 10% between the winner and the runner-up). For observations at election i , county j , and event month t : $Y_{ijt} = \alpha + \gamma_1 1\{\text{vote margin} > 0\}_{it} + \delta_0 [\text{vote margin}]_{it}^- + \delta_1 [\text{vote margin}]_{it}^+ + \varepsilon_{ijt}$ where “vote margin” is the percentage point vote margin between the Black candidate and the White candidate. Column (3) imposes a tighter 5% bandwidth. Columns (4)-(6) uses optimal bandwidths from Calonico et al. (2014), with linear, quadratic and cubic control functions and triangular kernel. Estimates for second- and third-order polynomial terms for the post-period – as well as all polynomial estimates for the pre-period – are not displayed. Columns (7)-(9) replicate columns (1), (2) and (4) with the “mandatory” sub-sample. Observations are weighted by the fraction of the geography (e.g. county) affected by the election. Standard errors are clustered by election $\times$ county.

	Post								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Black winner	0.025 (0.010)	0.025 (0.012)	0.017 (0.014)	0.021 (0.010)	0.023 (0.011)	0.026 (0.012)	0.025 (0.012)	0.035 (0.012)	0.024 (0.010)
Margin (winner)	-0.004 (0.001)	-0.001 (0.002)	0.002 (0.005)	-0.000 (0.001)	-0.001 (0.001)	-0.003 (0.002)	-0.003 (0.002)	-0.002 (0.002)	-0.000 (0.001)
Margin (loser)	0.001 (0.002)	-0.000 (0.002)	0.000 (0.006)	-0.000 (0.001)	0.000 (0.002)	-0.000 (0.002)	-0.001 (0.002)	-0.003 (0.002)	-0.000 (0.001)
Observations	23590	21462	14672	37008	47646	62914	18177	16481	32058
	Pre								
Black winner	0.003 (0.009)	0.008 (0.010)	-0.008 (0.012)	0.005 (0.009)	0.008 (0.010)	0.012 (0.010)	0.012 (0.012)	0.022 (0.014)	0.000 (0.011)
Observations	23437	22271	15083	37291	47875	62212	15319	14421	27604
Demo-adjusted		O	O	O	O	O		O	O
Only mandatory							O	O	O
Bandwidth	10	10	5	23.8	32	45.7	10	10	28.2
Polynomial	1	1	1	1	2	3	1	1	1
Triangular kernel				O	O	O			O

Table 5: Heterogeneity

The table display results from the regression-discontinuity design, with heterogeneous effects. The specification imposes a 10% bandwidth (i.e. regression sample only includes elections for which the vote margin is at most 10% between the winner and the runner-up). For observations at election i , county j , and event month t , the following regression specification is run and columns (1)-(3) report the two γ_1 on the Black-victory dummy and its interaction with another indicator for the county having above-median IAT in 2005, Black population share or income respectively:

$$Y_{ijt} = \{ \alpha + \gamma_1 1 \{ \text{vote margin} > 0 \}_{it} + \delta_0 [\text{vote margin}]_{it}^- + \delta_1 [\text{vote margin}]_{it}^+ \} \\ \times \left\{ 1, 1 \{ \text{county has high level of ...} \}_j \right\} + \eta_{ijt}$$

where “vote margin” is the percentage point vote margin between the Black candidate and the White candidate. The first indicator term denotes a dummy for elections in which the Black candidate won. The next two terms are linear controls separately for elections where the White candidate won (i.e. negative vote margin) and where the Black candidate won (i.e. positive vote margin). Observations are weighted by the fraction of the geography (e.g. county) affected by the election. In the last three columns, the RD regression is run for sub-samples, divided by office types.

	By county's			By office type		
	(1) IAT	(2) Black pop	(3) Income	(4) Mayor	(5) US House	(6) Other
Black winner	0.014 (0.017)	-0.001 (0.021)	0.057 (0.024)	0.060 (0.022)	0.115 (0.057)	0.014 (0.013)
× high level	0.018 (0.023)	0.029 (0.025)	-0.054 (0.027)			
Observations	21462	21462	21462	1574	1006	18882

Table 6: Economic gaps: RD and IV estimates

This table displays both the reduced-form RD results for economic gaps and the second-stage estimates of an instrumental variables regression, using the regression discontinuity (dummy for the Black candidate’s victory) as the instrument. The main specification imposes a 10% bandwidth (i.e. regression sample only includes elections for which the vote margin is at most 10% between the winner and the runner-up). For observations at election i , county j , and event year t , the following regression specification is run, and I report β_1 for different outcome variable Y_{ijt} :

$$Y_{ijt} = \beta_0 + \beta_1 \text{IAT}_{ijt} + \tilde{\delta}_0 [\text{vote margin}]_{it}^- + \tilde{\delta}_1 [\text{vote margin}]_{it}^+ + \varepsilon_{ijt}$$

$$\text{IAT}_{ijt} = \alpha + \gamma_1 1\{\text{vote margin} > 0\}_{it} + \delta_0 [\text{vote margin}]_{it}^- + \delta_1 [\text{vote margin}]_{it}^+ + \eta_{ijt}$$

where “vote margin” is the percentage point vote margin between the Black candidate and the White candidate. The first indicator term denotes a dummy for elections in which the Black candidate won. The next two terms are linear controls separately for elections where the White candidate won (i.e. negative vote margin) and where the Black candidate won (i.e. positive vote margin). Observations are weighted by the fraction of the county affected by the election. Standard errors in parentheses are clustered at the election $\times$ county level.

	Labor market		Mortgage		
	transition to unemployment	employment to pop	log originate to pop	rejection rate	
Reduced-form RD estimates					
Black winner	0.025 (0.011)	-0.050 (0.032)	-0.256 (0.111)	0.035 (0.015)	
Observations	4510	4376	4311	4166	
Adjusted R^2	0.004	0.023	0.024	0.027	
Second-stage IV estimates					
IAT (adj)	0.443 (0.237)	-1.234 (0.802)	-4.339 (2.534)	0.634 (0.349)	
Observations	3265	3188	3099	3015	
Transition-to-unemployment RD estimates by employee count					
	0-19	20-49	50-249	250-499	500+
Black win	0.030 (0.019)	-0.042 (0.035)	-0.003 (0.021)	-0.012 (0.017)	0.011 (0.014)