How Context Affects Politics: Essays on Causality and Measurement

How Context Affects Politics: Essays on Causality and Measurement

A dissertation presented

by

Melissa Lee Sands

to

The Department of Government

in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of
Political Science

Harvard University
Cambridge, Massachusetts

May 2017

Dissertation Advisor: Gary King Melissa Lee Sands

How Context Affects Politics: Essays on Causality and Measurement

Abstract

How do contextual features of the neighborhoods where we live, work, and play influence civic and political behavior? Micro-level forces, often overlooked in the study of broad social and political processes, shape our decision-making in myriad ways. I propose that who and what we encounter as we go about our everyday lives affects the policies we support, as well as how we vote and interact with government. This dissertation explores the effects of social and physical context through a series of empirical studies, beginning with the most micro of settings — individuals on a city street — and progressing to the more macro — voting precincts, neighborhoods, and municipalities.

The first study describes a placebo-controlled field experiment in which encounters between the well-off and the poor in affluent neighborhoods decrease the willingness of the former to support redistributive public policies. The second study uses geo-located pedestrian traffic counts from a network of live video camera feeds, overlaid with municipal hotline (311) data and crime incident reports, to demonstrate how the degree to which city-dwellers function as "eyes on the street" is highly context-dependent. The third study reveals how violent protests can be mobilizing events that shift local voter support in favor of public goods associated with the rioting group, and that this mobilization has enduring downstream effects. The final study leverages the end of legal, state-sponsored segregation in South Africa, along with natural physical geography, to show that sustained racial isolation increases white South Africans' likelihood of voting against historically black African or other non-white parties.

Contents

List of Figures

List of Tables

Acknowledgments

"Sink or swim" is how Harvard is described. If I didn't sink, it was because I was lucky enough to have mentors, advisors, friends, and family to keep me afloat. Though I swallowed plenty of water in the process, the damage would have been far greater were it not for the individuals and institutions named below.

First, my unrivaled committee.

Gary King pushed me harder than anyone ever has. For that, and for providing sage advice and much-needed cheerleading, for being available (no matter where he was in the world), and for always welcoming disagreement, I am extraordinarily grateful. I hope that ultimately I was able to write a dissertation that is, in his words, "worthy of Melissa."

Nobody has a mind like Dan Carpenter. I could always count on him to put forth some incredibly insightful suggestion or comment that led me to see my own ideas in a new light. In particular, I am eternally grateful for his support as an advocate and advisor during the roughest patches of my graduate career.

No one has had a greater impact on my graduate career than Ryan Enos. He is the kind of scholar, advisor, and mentor that I aspire to be. His contributions to this dissertation, and to how I think about human behavior, are too many to list. Simply put, Ryan brought out the scientist in me. Though I never indicated as much, there were multiple points at which he kept me from switching careers or leaving graduate school altogether. Ryan, I cannot thank you enough for believing in me even when I failed to believe in myself, as well as for creating a community in which graduate students like me can thrive at Harvard.

In addition to my committee, Arthur Spirling was vital to my survival and success in graduate school. I doubt he realizes how much he helped me, both professionally and personally, especially during that rocky patch in my fourth year. I pledge to honor his mentorship — and that of Dan, Gary, and Ryan — by paying it forward.

A number of other Harvard faculty members have been generous with their time, including Claudine Gay, Jim Snyder, Steve Ansolobehere, Ken Shepsle, Matt Blackwell, Maya Sen, and Chase Harrison. Laura Donaldson, Lilia Halpen-Smith, and Thom Wall deserve my gratitude as well.

This dissertation would have not have been possible without the generous support of Harvard funders, including the Institute for Quantitative Social Science, the Center for American Political Studies, and the Pershing Square Venture Fund for Research on the Foundations of Human Behavior. Jason Novack and Martin Lagache at Placemeter generously shared data and feedback. A series of talented undergraduate Research Assistants, including Eliza DeCubellis, Cara Kennedy Cuomo, and Quinn Mulholland, were crucial to the first paper.

To all past, present, and future members of the Government Department Diversity Working Group — thank you for your crucially important service to your peers and to the discipline at large. May the powers that be never underestimate your value and contributions.

I am also grateful to the faculty and graduate students at UCLA who made my five month stay in 2014 enriching and enjoyable. I'm particularly thankful for the continued mentorship and friendship of Miriam Golden and David Sears.

I would not have embarked on this journey if it were not for several brilliant and extraordinarily generous individuals I met at the School of International and Public Affairs at Columbia who continue to shape my life in various ways. Matt Rae helped me fall in love with statistics and continues to inspire me with his infinite array of talents and his boundless munificence. Ryan Hübert facilitated my love of causal inference and guided me as I took the leap into the mysterious gauntlet of Ph.D. programs. Mike Ting was my first academic mentor; our long conversations while I was but a naïve MPA student motivated me to pursue research as a career. Ken Prewitt took me under his wing and shaped my broader ambitions; he is a model for how academics can influence public policy without compromising their integrity as scientists and researchers. I hope to one day make him proud.

The long slog of the last six years would have unbearable without the incredible friendship and generosity of countless individuals in and around Harvard and MIT, as well as in the discipline more broadly. To everyone who has helped to celebrate the variegated victories and milestones, to mourn the various losses and disappointments, and to fill our little Irving Street abode with laughter and joy, thank you. To everyone who commented on ideas, research designs or drafts, patiently sat through presentations and practice job talks, or listened to me rant about political science, politics, and the patriarchy, thank you. Although there are too many individuals to name here, Stephen Pettigrew, Anton Strezhnev, Dana Higgins, Michael Hankinson, Ariel White, Chris Celeya, and Chris Lucas deserve special acknowledgement. Many, many others have enriched my life both personally and professionally during my time at Harvard. They include Kris-Stella Trump, Alex Hertel-Fernandez, Michael Gill, Jesse Gubb, Tess Wise, Rakeen Mabud, Emily West, Leslie Finger, Jason Sclar, Mayya Komisarchik, Elena Laudet, Justin de Benedictis-Kessner, Sole Artiz, Peter Bucchianeri, Shelby Grossman, Dominika Kruszewska, Riley Carney, Amy Uden, Shiro Kuriwaki, Dan Moskowitz, Anthony Fowler, Molly Roberts, Hye Young You, Bernard Fraga, Max Palmer, Andy Hall, Ben Schneer, Kara Ross Camarena, Aaron Kaufman, Leah Stokes, Chad Hazlett, Dean Knox, Horacio Larreguy, Ben Newman, Amy Catalinac, Dena Enos, and Josh Kertzer.

Outside of the political science bubble, Liz Lazar, Jenni Gallico, and Andy Bernstein — and their amazing partners, Marc, Max, and Jane — kept me sane these last six years. Their friendship means the world to me. Though I was not always the best at returning phone calls, I would replay their voice messages just to hear voices that would bring me home when I felt so far away. Thank you for not giving up on me all these years.

Surviving — and ultimately succeeding — in graduate school takes stubbornness, courage, and grit. I learned these traits from my parents, Rona and Jeffrey Sands. My mother, more than anyone else, has shaped the person I am today. She taught me to ask questions, to love learning, to pursue my dreams, to advocate for myself, and to stand up for what is right. The values and integrity she instilled in me are reflected in my goals and aspirations. *Tzedek, tzedek tirdof.* From my father, I learned the true meaning of unconditional love. He showed me how to use humor to defuse conflict and cope with challenges; to be tenacious, ambitious and entrepreneurial; and to balance a sharp skepticism with a willingness to take risks. He taught me not to believe the hype: success and brilliance don't require a fancy degree.

My twin brother, Justin, blows my mind with his profound serenity and ability to acquire deep knowledge of whatever he sets his formidable brainpower upon. I am also grateful to my incredible aunts, uncles, and cousins — some of whom I got to know only recently, but are no less impactful on my life. Though I cannot name them all here I want to acknowledge Annice Samskin, who has been like an older sister to me and whom I count among the most important people in my life. My newly-acquired South African family — Elizabeth and Raphael, Julia and Nick, Chris and Gina — have been endlessly generous and accepting, despite my many strange American quirks.

At each academic milestone I think of my grandparents: Roslyn and Morris, and Goldie and Leo. If they were alive today, I suspect they would be the proudest of all. None of them had the opportunity to attend college; three out of the four had to drop out of high school to help support their families. The sacrifices they made were so that their children, and their children's children, could have a better life. Though they surely had lofty dreams for their offspring, I doubt any of them could have imagined that one day their granddaughter would earn a doctorate from Harvard. This dissertation is dedicated to them.

Your spouse is the one family member you get to pick, and indeed the best decision I ever made was to marry Daniel. His patience, generosity, thoughtfulness, affection, and sense of humor are unparalleled. After nearly four years of marriage, I am still in awe of him, and, especially, in his willingness to put up with me day after day. His imprint is on every page of this dissertation; every idea, every project, and everything written here has benefited from his sharp eye and exceptional mind. He has assisted me, without compensation or complaint, in the field and in the office, and accompanied me on countless "work" trips at my request. He is the rare husband who puts his wife's career first, a sacrifice that means the world to me. I am, unequivocally, much better off for this two-body blessing.

Introduction

Imagine identical twins, raised under the same roof in Utah, their upbringing and experiences indistinguishable through young adulthood. Their social and political identities —as conservative, poor, white men, and self-avowed Republicans —crystallizes in those formative years. Now imagine that an omnipotent social scientist flips a coin to determine what happens next: one twin is sent to live in the affluent and highly-educated liberal bastion of Brookline, Massachusetts, and the other to rural, conservative Lancaster county, Pennsylvania. How does the result of the coin toss affect the twins' political behavior? For which party's presidential candidate will the two men vote in future elections?

This is the thought experiment that, in the pages of *Political Geography* in 1996, opened a debate between scholars of Political Science and Geography about the role of context in explaining political behavior (Agnew 1996a,b; Flint 1996; King 1996). Political scientist Gary King offered this hypothetical, arguing that the twins will continue to support the Republican candidate in presidential elections after they relocate. Context, he claims, should not count. Geographer John Agnew, meanwhile, doubled down on the importance of context. He asserts that the political scientist's conception of context is too narrow, too local. Understanding political behavior, he writes, requires studying the global social-geographical forces in which individuals are situated.

King and Agnew, as others have noted since, talk past one another in this now decades-old exchange. One key reason is their fundamentally different definitions of what constitutes a *contextual effect*. King's thought experiment makes sense if we wish to make a causal inference: what is the treatment effect of random assignment to some context (say, Brookline), versus the counterfactual assignment to another context (say, Lancaster). Agnew's macro conception of context focuses instead on situating individuals within geographical levels from which they and their behavior are necessarily inseparable; in this construction, envisioning a counterfactual like King's is not only irrelevant but nonsensical.

Much of political science is concerned with this macro conception of context effects: how do big economic and geopolitical events or institutions influence political or social outcomes? By contrast this dissertation embraces King's more neglected characterization of contextual effects as a micro-level phenomenon: how does who and what we encounter as we go about our everyday lives affect the policies we support, and how we vote and interact with government?

0.0.1 Taking Geography out of the Residual

On the face of it, King's assertion that our hypothetical twins will ultimately vote for the Republican candidate is a fairly uncontroversial statement. Political science teaches that, in the U.S. at least, party affiliation is a social identity (Green, Palmquist, and Schickler 2002) which individuals more or less inherit from their parents. Party identification tends to be stable over time within a given individual (Campbell et al. 1960). In the short run at least, political scientists tend to think that features of who you are — racial and ethnic background, your educational attainment, your religious affiliation or lack thereof — are more important than where you are in space. This is not to deny the explanatory power of geographic variation for Americans' political attitudes; the partisan divides of urban vs. rural, coastal versus land-locked, do follow a spatial pattern familiar to even the most casual observer of politics. Yet observable variation over space does not inherently imply that space or context exercises any causal effect. Residential sorting in particular is thought to be a pervasive phenomenon (Schelling 1971b,a) to the extent that individuals select into neighborhoods populated by people like them, the relationship between political attitudes and geographic context is largely confounded. Even when sorting is not thought to be a primary concern, such as in V.O. Key's 1949 classic description of how racial context influences political behavior in the U.S. South, reliance on observational data is fundamentally limiting. In the absence of a truly comprehensive set of individual-level variables that explain behavior, researchers can easily misattribute the "effect" of individual characteristics to context.

For these reasons it is easy to agree with King's assessment of geography as the residual category; it has explanatory power only to the extent that researchers lack the "right" individual-level variables and fail to fully understand human behavior. Indeed, in 1996 political scientists had not (yet) convincingly demonstrated that causal context effects exist. As King observes, "...in the numerous studies of voting and other political behavior, the existing evidence is fairly thin for any argument about major contextual effects." (King 1996)

Yet in the intervening twenty years researchers have designed and implemented studies that explicitly distinguish context effects from geographic confounders. Thanks in large part to methodological advances in political science and the broader social sciences, these constitute a growing body of evidence in support of racial, social, and physical contextual effects. They include natural experiments (Berger, Meredith, and Wheeler 2008a; Ananat and Washington 2009; Hopkins 2010; Ananat 2011; Acharya, Blackwell, and Sen 2016; DeCelles and Norton 2016; Enos 2016; Trounstine 2015), large-scale policy experiments, such as Moving to Opportunity (Katz, Kling, and Liebman 2001; Kling, Liebman, and Katz 2007; Gay 2012; Chetty, Hendren, and Katz 2016), and field experiments (Enos 2014). Though these types of studies are particularly challenging to conduct — truly random variation in context rarely occurs naturally in the real world, and is difficult to convincingly, cost-effectively, and ethically execute in a field experiment — they provide some of the best evidence in favor of the claim that context counts in real and important ways. Complementing these developments in studying causal effects, technological and policy developments have augmented researchers' capacity to measure attributes of context. For example, fine-grained geospatial data is increasingly available through 'open government' initiatives, allowing researchers to study the most local of conditions.

The chapters that follow build on and extend this relatively young line of research, investigating the political effects of social and physical context through a series of empirical studies. Each contributes to our understanding of how contextual features of the neighborhoods where we live, work, and play influence civic and political behavior. We proceed from the most micro context to the more macro: from individuals to city blocks, to precincts and neighborhoods, to larger voting districts and municipalities.

0.0.2 Study 1: Exposure to inequality affects support for redistribution

The first study approximates the twins thought experiment by randomizing individuals' exposure to contexts of inequality that have been experimentally induced in the real world. Micro-level forces — small-scale interactions and contextual features — that shape human behavior have been largely ignored in studies of inequality. Yet social psychologists have found that individual experiences, such as those that shape one's frame of reference, are crucial to decision-making. As the world's population grows more urban, encounters between members of different socioeconomic groups occur with greater frequency. Study 1 reveals that these interactions shape people's revealed preferences for redistribution.

In a pre-registered placebo-controlled field experiment, I show that encounters between the well-off and the poor in affluent neighborhoods decrease the willingness of the former to support redistributive public policies. I randomly assign micro-environments of inequality by placing "poor" confederates in affluent public spaces. This overcomes the selection bias that confounds most research on context effects. Passersby were asked to sign a petition calling for greater redistribution through a Millionaire's Tax, or, to establish a baseline, a placebo petition unrelated to inequality. Confederate race was also randomized, allowing me to demonstrate that race and poverty cues interact to sway people's decisions about which policies to support. Results from nearly 2,600 solicitations show that in an everyday setting exposure to inequality negatively affects willingness to publicly support a redistributive policy. The effect is principally due to encounters between those who are similar in terms of race and gender. I propose that social comparisons drive the observed heterogeneous effects.

That under certain conditions inequality begets inequality has fundamental implications for policymakers. The findings also inform our scientific understanding of the effects of economic and racial inequality, set against a backdrop of segregation. Crucially, certain manifestations of inequality trigger reactions in some groups of people but perhaps not in others. Methodologically, the paper provides a new experimental design for exploring these effects in the real world.

0.0.3 Study 2: 'Eyes' on the street: What public camera feed data can teach us about civic and political behavior

While the experiment in Study 1 directly manipulates features of context, Study 2's emphasis is on measuring aspects of context. I use geo-located pedestrian traffic counts from a network of live video camera feeds to study several social phenomena. Data on foot traffic, singular in its precision and temporal coverage, comes from a private company that uses a network of live video camera feeds and computer vision technology to calculate traffic levels in real time. After discussing the data structure and limitations, I present several social science applications that describe the relationship between space and 'civic-ness' in a highly walkable urban setting.

First I show that these data can be used to study how crowds develop, and then disperse, during large political protests. Next, utilizing publicly available geospatial data from New York City, I test Jane Jacobs' 1961 classical theory about what makes city sidewalks safe and vibrant. To do this, I use spatial kriging techniques to interpolate hourly foot traffic in Manhattan, which is then overlaid with geo-located New York City 311 reports and crime incident reports. I demonstrate that higher foot traffic, and therefore more "eyes on the street", does not necessarily reduce crime or encourage people to report non-emergency issues (via municipal hotlines). Rather, the degree to which urban denizens function as "eyes" is highly context-dependent and non-linearly related to foot traffic.

0.0.4 Study 3: Can violent protest change local policy support? Evidence from the aftermath of the 1992 Los Angeles riot (with Ryan D. Enos and Aaron R. Kaufman)

While context may influence behavior directly as shown in Study 1, it also matters in other ways. For example, major political events such as violent protests can have important local consequences that tend to be ignored in political science research. While a broad focus is valuable, persistent attention to more general phenomena may conceal differences between local and distal effects. A riot represents a type of shock to community members' context; those proximate are most likely to be materially and psychologically affected, and thus most likely to be mobilized in the aftermath. If we wish to liberate political behavior from the proverbial black box, and to understand the forces that sway local politics and policy, it behooves us to examine the localized contextual effects of such events, even as they gain national attention.

To that end, Study 3 examines the effect of the 1992 Los Angeles riot —one of the most high-profile events of political violence in recent American history— on policy support and partisan voter registration among African American and white L.A. residents. Using precinct-level ballot initiative results and geocoded voter registration records – both real behavioral outcomes from before and after the riot – we document a liberal shift in observed political behavior following the riot. To isolate the causal effect of the violent protest, which occurred just prior to an election, we use the change in the difference between voter support for public education and higher education as a proxy for attitude shifts. We find that voters in the L.A. basin become more supportive of providing a public good closely associated with the rioting group, local public schools. To investigate the source of this shift, we examine survey data and voter registration records for those who registered just before and just after the riot. Ultimately we attribute the vote swing to increased mobilization of both African American and white voters.

The findings also show that these local contextual effects are anything but ephemeral.

To examine the downstream of effects of registration, registrants on the 1992 voter file were matched to those on the 2005 voter file. High stability in both participation and in partisanship among those who registered after the riot emerges, suggesting long-term consequences of the riot: citizens, white and non-white alike, appear to have joined the party sympathetic to the demand of the rioters and not only remained in that party, but stayed active citizens. The finding that mobilization endures over a decade later speaks to potential durability of contextually dependent events.

0.0.5 Study 4: Segregation drives racial voting: New evidence from South Africa (with Daniel de Kadt)

In countries across the globe, one of the most pervasive features of how people are distributed across space is segregation, or the physical separation of groups along racial, ethnic, religious, or economic lines. Segregation can be the deliberate result of state or local policy, or it can result from residential sorting or selection into relatively homogeneous enclaves. Study 4 directs attention beyond the U.S. to the effect of local demographic shifts that occurred at the end of legal segregation in South Africa. We demonstrate that changes in local racial context precipitated by the end of Apartheid and the advent of inclusive democracy in 1994 affect white citizens' willingness to vote across racial lines. In particular, we show that racial isolation increases whites' likelihood of voting against historically black African or other non-white parties.

To do this new datasets are required, including for the first time high resolution census data from before the end of Apartheid. Several quasi-experimental research designs are used to make causal inferences, leveraging natural physical geography — the hills, valleys, and ridges that act as buffers against racial mixing — as an instrument for sustained racial isolation. To address ecological inference concerns and explore potential mechanisms, geo-referenced survey data is used to provide individual-level evidence consistent with our findings. As the U.S. continues to deal with its own legacy and perpetuation of segregation, the South African case sheds light on the broader consequences of racial isolation on voting behavior.

1 Exposure to inequality affects support for redistribution

1.1 Introduction

The distribution of wealth in many countries across the globe is highly skewed. In the United States the gap between the top one percent of earners and everyone else was wider in 2012 than any time since before the Great Depression (Piketty and Saez 2013), and is a pervasive social and political phenomenon (Bartels 2009). With this comes visible manifestations of inequality, which affect social interactions from cooperation (Nishi et al. 2015) to conflict (DeCelles and Norton 2016). Yet very little is known about how direct exposure to inequality in everyday settings – such as poverty and homelessness in relatively wealthy neighborhoods – shapes human behavior. Isolating this causal effect is challenging because of selective sorting. To overcome this, I experimentally manipulate exposure to inequality that occurs on sidewalks and street corners. Using a randomized placebo-controlled field experiment I show momentary passive "contact" with a poverty stricken person in an affluent public place can change people's willingness to actively support redistributive policies.[1]

Economic inequality is an abstract concept, difficult to concretely portray without the help of numbers, graphs, or words. Understanding the implications of exposure to inequality as a *personal experience*, rather than as an impersonal abstraction, requires the manipulation of micro-level contextual features. I evoke everyday inequality by placing poor individuals

[1] A slightly modified version of this paper was published as: Sands, Melissa L. 2017. "Exposure to inequality affects support for redistribution." *Proceedings of the National Academy of Sciences* 114(4): 663–668. Replication data are available on the Harvard Dataverse at Sands (2017b).

in a place of affluence.[2] Those with the means to live and shop in such neighborhoods are more likely than the average citizen to participate in politics and donate to political causes, and thus wield a disproportionate influence over politics and policy (Bartels 2009).

The experimental intervention created a micro-environment of inequality that was both highly realistic and carefully controlled by the researcher. In treatment conditions, passersby were subtly exposed to a professional actor (confederate) portraying an impoverished person through both their clothing and body language. In the control condition, the same confederate would portray an affluent individual, dressed business casual and showing appropriately different body language.

Since the 1960s racial attitudes have dominated the American public's thinking about poverty (Gilens 1999). Political scientists have shown that simply priming race in a survey can lower White respondents' support for welfare policies, e.g.,(Peffley, Hurwitz, and Sniderman 1997; Gilens 1996). Yet social psychology demonstrates that individuals define themselves in relation to similar others in their environment, such as those who share their race and gender (Wood 1989). This phenomenon might be particularly relevant in an experiential setting. To assess whether predominantly White and well-off pedestrians respond differently to inequality that appears both socioeconomic and racial, the race of the confederate was also randomized.

Passersby became subjects in the experiment as they were approached by a petitioner, stationed within twenty feet of the confederate. The policy on the petition was randomized between two texts, the first of which expressed support for a "Millionaires' Tax" to increase taxation on incomes over $1 million, redistributing the money. To establish a baseline, the other half of subjects were asked to sign a placebo policy petition unrelated to inequality and poverty, referring instead to reducing disposable plastic bags.

Results from 2,591 solicitations show that individuals are significantly *less likely* to support redistributive policies in the presence of a poor person in an affluent setting. The

[2] Using a separate, online experiment (detailed in the appendix), I show that subjects describe images of the poor in affluent settings as depicting inequality. Thus I use the terms "inequality" and "poverty in an affluent setting" interchangeably.

finding holds overall and when estimated as a difference-in-differences (net of the baseline response rate established by the placebo petition). The negative causal effect of direct exposure to inequality on support for redistribution appears to be driven largely by White men encountering a poor person *of their own race*.

This study makes several methodological and empirical contributions. First, I provide real-world field experimental evidence that direct exposure to inequality – something that occurs daily worldwide – influences support for redistribution. Theoretically, this implies that momentary contact between highly different socioeconomic groups may have the perverse effect of lowering wealthier people's support for redistributive policies. It also helps elucidate the contextual factors that shape support for redistributive policies, including how the presence of poverty in a place of affluence influences political preferences. Second, by randomizing the race of confederates as well as their socioeconomic status, this study experimentally disentangles the confounded relationship between race, poverty, and support for redistribution. I find that White subjects respond negatively to poor White confederates, but not to poor Black confederates.

A handful of studies have examined the effect of context or subtle environmental cues on behavior (Berger, Meredith, and Wheeler 2008a; DeCelles and Norton 2016) and attitudes (Enos 2014). While none have randomized encounters with real-world inequality, several recent studies examine the effect of perceptions of inequality on survey responses and laboratory behaviors (Côté, House, and Willer 2015; Trump 2017; Nishi et al. 2015; Newman, Johnston, and Lown 2015). In a series of laboratory experiments Nishi et al. (2015) show that visible wealth encourages greater inequality, compared to when wealth is hidden. Using survey experiments, Côté, House, and Willer (2015) show that higher-income individuals are less generous, relative to lower-income individuals, if they reside in a highly unequal state or when randomly assigned to view simulated data indicating inequality is relatively high in their state. Rather than presenting subjects with information about inequality, this study experimentally induces inequality in a person's natural, ordinary environment. Further, prior research on inequality effects has typically occurred in a relatively artificial context in which, among other concerns, outcomes are subject to experimenter effects (Rosenthal and Fode 1963; Schultz 1969). I measure revealed preferences through a real-world behavior, in a context in which experimenter effects are highly implausible. In contrast to previous studies of redistributive preferences that rely on costless expressions of opinion on a survey or relatively unnatural distributive games, this study examines the effect of exposure on a real political action in an authentic setting. Although signing a petition is less costly than some forms of political engagement, it is a meaningful, public form of advocacy (Carpenter and Moore 2014; Han 2016) and a potential gateway to future political participation (Parry, Smith, and Henry 2012; John et al. 2013; Lee and Hsieh 2013).

1.2 Research Design

The experiment took place between October 3rd and December 11th, 2015. Study sites were affluent, predominantly White, pedestrian-trafficked commercial areas around Boston, Massachusetts.[3] A fully-crossed factorial design was employed, depicted in table 1.1. The "poverty" and "affluence" conditions refer to the appearance of the confederates. In the poverty condition, confederates appeared unkempt, wearing extremely shabby clothes and communicating poverty via their posture and body language (see figures 1.2 and 1.4). In the affluent condition, the same confederates were well-dressed and behaved as if waiting patiently to meet a companion for a nice meal (see figures 1.1 and 1.3). Confederates, positioned within twenty feet of the petitioner, were instructed to remain standing so as to be visible to pedestrians, and to not, in any way, appear to be begging or panhandling.[4] They were told to refrain from initiating conversation with pedestrians, petitioners, or anyone else while in sight of their assigned location.

On any given day, two male confederates of roughly the same age, one non-Hispanic

[3]Neighborhoods included Brookline, Back Bay, and Beacon Hill, and were chosen to limit the heterogeneity of the subject pool in order to ensure sufficient statistical power.

[4] Confederates were instructed not to panhandle, and not to appear mentally ill or intoxicated, to avoid generating aversion, scorn, or disgust. All six confederates were trained and monitored closely by the researcher.

Table 1.1: Experimental Design

Petition	Actor	**Poverty**	**Affluence**
Millionaire's Tax: Increasing taxes on earnings over $1 million, redistributing the money.	**White**	Treatment I	Control I
	Black	Treatment II	Control II
Plastic Bags: Reducing the use of plastic bags (unrelated to inequality or poverty).	**White**	Placebo I	Control III
	Black	Placebo II	Control IV

Note: This fully-crossed 2x4 factorial design describes eight conditions corresponding to two different petitions, two confederate race categories, and two confederate socioeconomic conditions. Each of the eight quantities can be expressed as a response rate, i.e., number of signatures divided by number of individuals approached in that condition.

White and one Black, portrayed "poverty" and "affluence" at randomly scheduled times throughout the day. The petitioner rotated between the two petitions, a millionaire's tax[5] and reducing plastic bags. Two conditions for each of confederate appearance, confederate race, and petition generate eight total conditions.

The plastic bag (placebo) petition accounts for the possibility that pedestrians behave differently in the presence of confederates regardless of the content of the petition. Response rates might always be lower when the petitioner is in the vicinity of a poor person because pedestrians are simply more hesitant to stop (a "scare-off" effect). The placebo allows me to rule out this possibility by taking the "difference-in-differences" in response rates for any given day and location. This nets out scare-off effects, isolating strictly the effect of the poverty and race treatments on the probability of signing the millionaire's tax.

The set-up, depicted in figure 1.5, was highly realistic in several respects. First, both petitions pertained to relevant policy issues in the state. During this time, activists were

[5] The so-called "millionaire's tax", a proposal to raise taxes on high income earners, is a salient and popular redistributive policy proposal on the agenda in many U.S. states. In MA, the policy would alter the state constitution to impose an additional tax of 4 percent on individuals with annual incomes of $1 million or more.

Figure 1.1: “Affluent” confederate

Figure 1.2: “Poor” confederate

Figure 1.3: “Affluent” confederate

Figure 1.4: “Poor” confederate

circulating petitions to get the millionaire's tax on the state-wide ballot in Massachusetts.[6] It was also realistic in that it is not uncommon to see poor-looking individuals standing around or sitting in these neighborhoods, though particular spots on sidewalks were avoided if occupied by a panhandler or someone who appeared to be homeless.[7]

To buttress the claim that the experimental intervention approximates a salient way non-poor individuals experience inequality, I conducted an online survey experiment in which an independent sample responded to photographs of scenes like those staged in the field. Subjects described those images in words and phrases related to inequality in open-ended questions, and ranked "inequality" as the best word to describe them.[8] This experiment is detailed in the appendix.

Although the primary research question is about exposure to inequality, operationalized through the comparison between the poor and affluent treatments, the factorial design of this experiment allows for comparisons across both dimensions of the treatment – socioeconomic status and race – and the interaction between the two. Because of the high salience of race in the U.S., subjects' reactions to the poverty and affluence conditions may be moderated by the race of the confederate. Thus in addition to the comparisons between row cells described in table 1.1 (treatment I vs. control I; treatment II vs. control II) I also describe the relationships between column cells (treatment I vs. treatment II; control I vs. control

---

[6] The campaign to place the millionaire's tax constitutional amendment on the ballot collected over 157,000 signatures, more than twice the 64,750 signatures needed to advance the policy to the legislative approval stage before it can be placed on the ballot in November, 2018. According to Raise Up Massachusetts, the advocacy coalition that sponsored the campaign, "The constitutional amendment would create an additional tax of four percentage points on annual income above $1 million. The new revenue generated by the tax...could only be spent on quality public education, affordable public colleges and universities, and the repair and maintenance of roads, bridges, and public transportation." `http://raiseupma.org/constitutional-amendment-campaign/`

[7] Importantly, in approximately thirty hours of petitioning, nobody expressed suspicion regarding the confederates nor the petitioner. The confederates changed their clothing for the different roles in a car with tinted windows that was parked on a side street, away from commercial strips where the petitioning took place. They were instructed not to interact with any member of the research team while in the vicinity of the petitioning location. Confederates walked to and from the car along back streets to avoid attention when not in character.

[8]Online subjects used terms of disparity and inequality to describe the images. When asked to rank words that describe the image, a majority ranked "inequality" as most aptly capturing photos of poor people in wealthy environments.

II).

The design was approved by Harvard University Committee on the Use of Human Subjects before the experiment or piloting began (Protocol IRB15-0930). A waiver of the consent process was obtained.

Figure 1.5: Experiment in Progress. The petitioner (**A**) approaches a subject (**C**) after the subject passes the confederate (**B**).

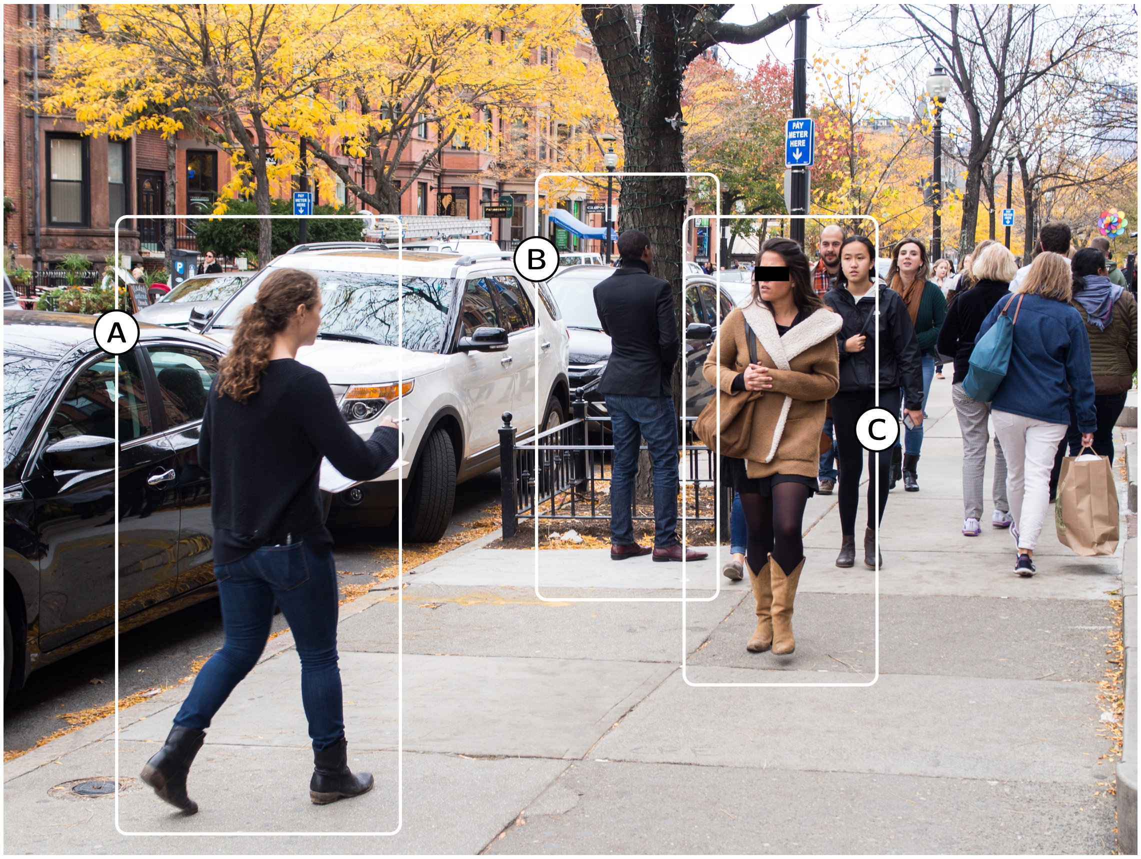

1.2.1 Randomization

Passersby were exposed to one of the eight conditions depending on when they happened to walk past the confederate in the direction of the petitioner. The experiment was complete block-randomized by day such that all eight conditions in table 1.1 occurred once each day,[9] rotating every thirty minutes. The starting condition and starting petition were randomized each day, and the confederates rotated based on these starting conditions. In total, there are 74 date-time clusters across 15 days. Average treatment effects are estimated at the cluster level as well as at the individual level with date-time clustered standard errors.

[9] On some days, only one confederate was available. When this occurred, all four conditions for that race were used. As a consequence, the treatment effect of race estimated in a model that includes day fixed effects is identified only by days on which both confederates were present. See appendix tables S3 and S4.

1.2.2 Petitioning

Measurement was conducted via response rates to a petition. The petitioner, a young, White, well-dressed student with clipboard and pen, counted pedestrians as they passed the confederate.[10] The petitioner approached every third adult, asking "Would you sign this petition in support of [the Millionaire's Tax / reducing the use of wasteful plastic bags]?" If questioned further about the nature of the policy, the petitioner responded with a brief, scripted reply, essentially paraphrasing the language on the petition.[11] If the subject asked to see the petition or agreed to sign it, he or she was handed or shown the clipboard. The exposed sheet displayed the text of the petition and a set of columns for printed name, signature, and address.

1.2.3 Covariates

Petitioners unobtrusively recorded their "best guess" at the gender, age, and race/ethnicity of each person approached, including those who did and did not sign. These covariates include the gender (man or woman), age (18-35, 35-65, 65+), and race/ethnicity (White, Black, Asian, Hispanic) of each individual approached. Table S1 in the appendix shows balance in the covariates across the affluence and poverty conditions, and table S2 shows that the eight treatment conditions predict covariates at a little worse than chance. Petitioners also kept track of unsolicited signatures and a set of other responses such as ignoring the petitioner. The tracking sheet is figure S5 in the appendix.

Based on the petitioners' recordings subjects were 82.8% White, 9.3% Asian, 3.9% Black,

[10] Petitioners were made aware of the presence of the confederate, and confederates and petitioners were introduced and familiar with one another. Petitioners were not given any information about the researchers' hypotheses, or their proposed analysis strategy, which was gated at EGAP. The decision to inform the petitioner of the confederates' presence was taken so as to minimize the possibility that the petitioners might behave differently in the varying conditions.

[11] Petitioners were given a short script (see appendix) to follow and were told not to deviate from it. They were also instructed not to try to persuade anyone to sign the petition through debate, or to engage in political discussions of any sort. In short, they "played dumb" in order to stick to the script. This ensured that their behavior remained consistent across experimental conditions. The four petitioners included female and male undergraduate students. There are no systematic differences in the results between petitioners.

2.5% Latino, and 1.5% unknown race. The sample comprised 57% women and 43% men, with 89% estimated to be between the ages of 18 and 65.[12] To estimate the socioeconomic status of those who signed a petition, addresses provided when signing were linked to zip code-level characteristics from the American Community Survey. Massachusetts residents, who comprise almost all of the signers, hail from high income zip codes, even by MA standards; the median zip code has a median household income of $79,290.[13]

In total, 2,591 solicitations were made across 74 time-date clusters. Each cluster represents a 30-minute period in which passersby were exposed to one of the eight conditions described in the cells in table 1.1. To account for dependence between observations within each 30-minute period, average treatment effects are calculated at the cluster level by regressing the response rate in each cluster on a set of binary indicators for the treatment conditions.

1.3 Results and Discussion

Cluster-level treatment effects are presented in table 1.2.[14] Based on column 1, subjects are found to be 4.4 percentage points *less* likely to support the redistributive policy in the presence of a poor person ($p<0.10$), pooling across confederate race conditions. The specification estimated in column 2 allows the treatment effect to vary by confederate race.

[12]Petitioners estimated 45% of the sample to be age 18-35, 44% age 35-65, and 11% over 65. Because guessing someone's age by sight alone is difficult, subjects are not analyzed separately by age group.

[13]This includes all MA zip codes that could be identified from the address information provided by the signer (n = 276), and is calculated from the 2010-2014 American Community Survey 5-year estimates of median household income in the past 12 months, in 2014 inflation-adjusted dollars. Although self-selected, petition-signers tend to come from neighborhoods where the median household income is well above that of the U.S.($53,657) and MA ($69,160).

[14]Millionaire's tax clusters include 1,335 subjects (solicitations), while plastic bag clusters include 1,256. See appendix table S3 for probit regression analyses, conducted at the individual level, with standard errors adjusted for 74 day-time clusters. Substantively, the main difference between individual and cluster level estimates is that the former are, essentially, weighted by the number of passersby in each cluster. The results are robust to specifications with and without day fixed effects and with and without covariate adjustment. They are also essentially the same when using logit regression or a linear OLS probability model. Given that the probability of signing any petition is on average below 0.25, the non-linear models are preferred for the individual-level estimates.

Here the estimated coefficient on *poor actor* nearly doubles in magnitude when accounting for the interaction between the poverty and race treatments. Exposure to White poverty in an affluent setting decreases support for a redistributive policy by 8.2 percentage points ($p<0.05$), a substantively and statistically significant decline.[15]

Meanwhile, subjects are *not* significantly more or less likely to support reducing the use of plastic bags under any of the conditions, based on columns 3 and 4 of table 1.2. These coefficient estimates are relatively small in magnitude and are statistically indistinguishable from zero.[16] Appendix table S6 shows estimates of the "difference-in-differences" between the plastic bag (placebo) petition and the millionaire's tax petition at the cluster level. The findings are highly similar; subjects are 8 percentage points *less* likely to support the redistributive policy in the presence of a poor person, net of baseline response rates ($p<0.05$), and this effect is driven exclusively by the poor White condition.

Figure 1.6 plots the difference in means, along with 95% and 80% confidence intervals, for several comparisons of interest in support for the millionaire's tax. Instead of estimating treatment effects at the cluster level, as in table 1.2, point estimates come from individual level data, in order to facilitate comparison with figure 1.7. Because each of the clusters is roughly the same size, individual-level estimates are nearly identical to those estimated at the cluster level.[17] As in column 2 of table 1.2, the difference in the poor Black versus affluent Black conditions is indistinguishable from zero, while the poor White versus affluent White comparison reflects the largest gap in response rates. Support for the millionaire's tax in the presence of an affluent Black confederate is lower than in the presence of an affluent White confederate, though the effect is not statistically significant at conventional levels.

[15] The response rate on the millionaire's tax petition declines from 14% in the affluent White condition to 6% in the poor White condition, a decrease of roughly 57%.

[16] Also see figure S1 in the appendix. The individual-level results, reported in tables S3 and S4, suggest that subjects are marginally less likely to stop to sign either petition in the presence of a Black confederate ($p<0.10$).

[17] Confidence intervals adjusted using a block bootstrap to account for dependence between observations within clusters are identical to unadjusted standard errors, suggesting there is no clustering in the individual errors.

Table 1.2: Cluster-level treatment effects

	Signed petition			
	Millionaire's tax		Plastic bags (placebo)	
	(1)	(2)	(3)	(4)
Poor actor	−0.044*	−0.082**	0.036	0.027
	(0.024)	(0.035)	(0.030)	(0.047)
Black actor		−0.051		−0.040
		(0.032)		(0.044)
Poor actor x Black actor		0.070		0.013
		(0.047)		(0.061)
Constant	0.114***	0.142***	0.180***	0.205***
	(0.016)	(0.024)	(0.021)	(0.034)
Clusters (observations)	38	38	36	36
Individuals	1335	1335	1256	1256
Residual Std. Error	0.073 (df = 36)	0.072 (df = 34)	0.089 (df = 34)	0.090 (df = 32)

Note: *p<0.1; **p<0.05; ***p<0.01

Cluster-level treatment effects from OLS regression. The first two columns show results for the millionaire's tax petition and the second two for the placebo petition. Columns 1 and 3 show the main effect of the poor actor treatment, pooling across Black and White actors. Columns 2 and 4 include the interaction between the race and poverty treatments.

Because this effect similarly appears among subjects asked to sign the placebo petition, it may be evidence of a scare-off effect caused by the hesitancy of subjects to stop for any cause when a Black person is nearby. The final comparison in figure 1.6 suggests that subjects are marginally less likely to support redistribution when the petitioner is adjacent to a poor White person rather than a poor Black person, though the difference is not distinguishable from zero.

Figure 1.6: Individual-Level Treatment Effects for Millionaire's Tax Petition (n = 1335).

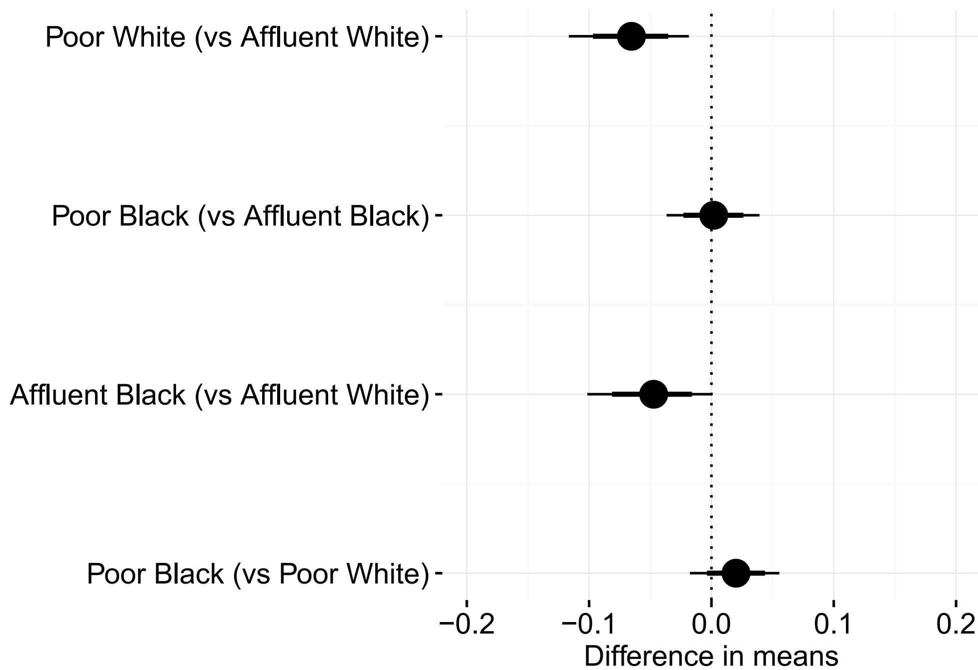

Note: Plot of the difference in means for all subjects on the millionaire's tax petition. Confidence intervals come from a block-bootstrap procedure to account for dependence between observations within clusters. Thin bars are 95% confidence intervals for two-sided t-tests on the difference in means; thicker bars represent 80% confidence intervals.

At the same time, the aggregate level results mask subgroup heterogeneity. Using the individual-level data recorded surreptitiously by petitioners, I examine non-Hispanic White subjects, who comprise over 80% of the sample, separately. Figure 1.7 shows the difference in mean response rates for White subjects (n = 1092), then for White men (n = 452) and White women (n = 640), on the millionaire's tax petition. Black subjects are not analyzed separately because there are so few of them (n = 62); they comprise less than 5% of the millionaire's tax sample.

Figure 1.7 shows the effect of inequality described above as primarily driven by White men, whose support for the millionaire's tax falls in the presence of a poor White confederate. As above, inequality due to a poor Black confederate barely registers for any subgroup. White men in particular are less likely to stop in the presence of an affluent Black confederate compared to an affluent White confederate, a "scare-off" that persists for both petitions (see placebo plots S1 and S2). Finally, though not statistically significant at conventional levels, White men are marginally more likely to support the tax in proximity to a poor Black person, compared to a poor White person, which further reflects their distaste for redistribution when a poor White man stands nearby.

These results reveal that visible inequality can be a pertinent feature of an individual's environment when asked to support redistributive policies. Why does the salience of everyday inequality *decrease* support for redistribution, particularly in some types of people?

One possible explanation is rooted in social psychology's "belief in a just world" hypothesis, which posits that affluent individuals justify their privileged positions by viewing the current distribution of resources as just and fair (Lerner 1980). As passersby in an upscale pedestrian thoroughfare, subjects may have reacted to a visual reminder of poverty by "doubling down" on their high economic standing. However, if inequality makes salient beliefs that the world is just, then subjects' support for redistribution should decline in the company of Black and White confederates alike. This hypothesis appears insufficient to account for the heterogeneous effects by race described above.

Another explanation is that visible poverty reminds people of their socioeconomic status relative to others. Individuals may adjust or benchmark perceptions of their own economic standing in relation to those in their environment (Merton 1957; Festinger 1954). Exposure to inequality may have induced subjects to switch frames of mind from considering themselves as the beneficiary of the tax to potentially being subject to taxation themselves (Cavaillé and Trump 2015); in other words, they "feel like millionaires" and thus balk at the

Figure 1.7: Treatment Effect on White Subjects for Millionaire's Tax Petition

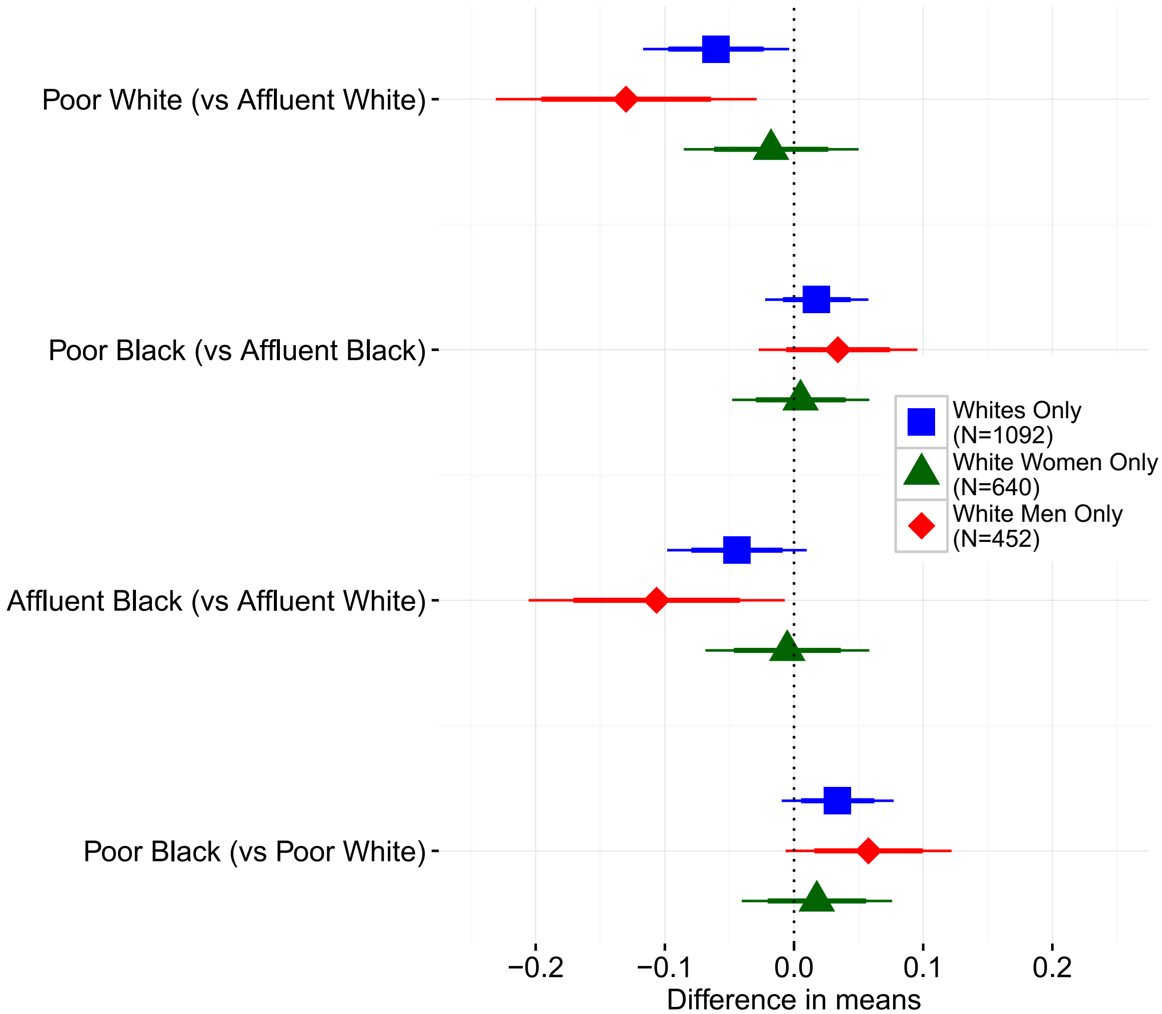

Note: Plot of the difference in means for White subjects, separated by men and women, in millionaire's tax petition responses. Thin bars are 95% confidence intervals for two-sided t-tests on the difference in means; thicker bars are 80% confidence intervals.

chance to support a policy that might harm them.[18]

The idea that individuals engage in "status benchmarking" against others is related to the social affinity hypothesis articulated by Lupu and Pontusson (2011), which describes how voters' preferences for redistribution should change with the perceived distance between high, middle, and low income groups. This explanation is also consistent with Condon and Wichowsky (2016), who find that reminding wealthy individuals of their relatively high social status can diminish support for redistributive policies.

Moreover, "status benchmarking" should occur most often when the congruence between subject and confederate characteristics is high; social comparison theory, well established in the social psychology literature, describes how individuals compare themselves to others with similar characteristics (Festinger 1954), such as race and gender (Campbell and Tesser 1985; Miller 1984, 1982).[19]

A related explanation is that the sight of a poor White person reduces support for redistribution among White subjects because the latter can more easily place blame on the poor individual, rather than society, for their misfortune. Individuals hold beliefs about what it takes to be successful *for members of their own group.* They learn from their own experiences with social mobility and develop attitudes about redistributive taxation based on those experiences (Piketty 1995, 1998). This leads them to believe that *effort*, rather than structural factors beyond one's control, predicts economic success; tangible reminders of inequality are seen through this lens, inducing them to reject redistributive policies. If this mechanism is at work, affluent White men believe the status of poor White men is due to lack of effort rather than structural factors, while less information is available to them about the effort exerted by poor Black men. This may lead White subjects to "punish" their poor White counterparts by rejecting the millionaire's tax. Indeed, individualized attributions of poverty have been shown to lower support for redistribution and welfare

[18] Petitioners reported that one reason people gave for not signing the millionaire's tax petition was that they were, in fact, millionaires, or they think that they will one day be millionaires.

[19] This occurs even when the basis of similarity is irrelevant to what is under evaluation. See (Wood 1989) for a review of the social comparison theory literature.

policies, as compared to structural or systemic attributions of poverty (Williamson 1974; Alston and Dean 1972; Iyengar 1990).

A more definitive test of the latter mechanisms would involve estimating treatment effects for affluent Black subjects exposed to Black or White poverty, as well as randomizing the gender of the confederates, both of which are outside the scope of this study. Still, the evidence presented here is consistent with the notion that subjects systematically respond most strongly to poor confederates that share their own race and gender. An individual who belongs to one's demographic in-group but is an economic outsider is a reminder of one's own status *within the relevant demographic group.* Thus we would expect people to reduce support for redistribution when visible inequality involves an individual with whom they share certain stable attributes.

1.4 Conclusion

This study uses a randomized placebo-controlled field experiment to establish the causal effect of exposure to inequality on support for redistribution. In doing so, it contributes to and extends several lines of prior scientific inquiry.

First it demonstrates the causal effect of encounters between groups of vastly different socio-economic status. This study also begins to disentangle the independent effects of race, gender and socio-economic status on political behaviors. Further, it is one of the first studies to examine the effect of exposure on an *actual political behavior*, in this case petition-signing. As such, our understanding of petition-signing as a political behavior is also developed – this behavior does not occur in a vacuum, though the context in which it takes place has been largely ignored.

Our understanding of the relationship between inequality and redistribution *at the individual level* is advanced. Much has been written on this topic at the national level, primarily claiming that rising inequality induces low-income voters to demand redistribution (Meltzer and Richard 1981; Roberts 1977; Romer 1975).This study contributes to an emerging experimental literature that examines this claim at a micro level, and is the first experiment to show that inequality in a real-world context may discourage affluent citizens from actively supporting redistributive policies.

The observation that momentary exposure influences behavior is suggestive of much broader consequences of repeated encounters between the wealthy and the poor. This type of "contact" is commonplace in cities, where population growth is occurring more rapidly than in suburban or rural areas (Cohen, Hatchard, and Wilson 2015). The prevalence of both urban growth and inequality in the U.S. means that understanding the behavioral implications of exposure to racial and economic outgroups is increasingly important.

The findings presented here suggest that the presence of poverty, particularly in a place of affluence, may decrease support for policies aimed at alleviating those conditions, a worrisome conclusion given that the general population increasingly resides in urban environments where contact with low-income individuals is likely. Homelessness, perhaps the most visible manifestation of rising urban poverty, may actually perversely discourage citizens from favoring social safety nets.

Finally, the results have practical implications for advocates engaged in canvassing, and political communication more broadly. Voters' reactions are contingent on the setting in which canvassing takes place, and perhaps even on the imagery utilized in campaign materials. Charitable organizations, fundraisers, and policy advocates may wish to exercise caution in using visual depictions of poverty or inequality. Ultimately, changing minds may require cognizance of context.

2 'Eyes' on the street: What public camera feed data can teach us about civic and political behavior

"The city at last perfectly illustrates both the universal dilemma and the general solution, this riddle in steel and stone is at once the perfect target and the perfect demonstration, of nonviolence, of racial brotherhood, this lofty target scraping the skies and meeting the destroying planes halfway, home of all people and all nations, capital of everything, housing the deliberations by which the planes are to be stayed and their errand forestalled."

– E.B. White, *Here Is New York*, 1949

2.1 Introduction

Researchers across a multitude of disciplines have long noted that *place* matters. From sociology, political science, and economics to urban planning, environmental psychology, and public health, there is substantial evidence that physical and social settings influence human behavior (Goffman 1971; Wolff 1973; Frank and Engelke 2001; Oliver 2001; Leyden 2003; Ananat 2011; Enos 2016). Cities in particular have been the subject of such inquiry. Yet we know remarkably little about how people move about in urban environments and how civic life plays out in densely populated areas.[1]

Such research remains nascent in part because the micro-level data required to study

[1]Replication data and materials are available on the Harvard Dataverse at Sands (2017c).

these processes is rare, and where available, difficult to use. Fortunately, technological advances – including those that have facilitated the ubiquity of cameras in public spaces, and the capability to process large streams of data – have fueled our ability to accurately measure, and thus systematically study, public citizen behavior more than ever before, most often in urban settings.

Cities are a particularly fruitful setting for the study of such political processes. They are unique in their density and forced concentration of individuals from diverse walks of life. *Other people* are a constant feature of the urban environment; as individuals interact with urban spaces they inherently interact with one another. This paper studies the quintessential example of an urban metropolis: New York City, where the island of Manhattan, a mere twenty-three square miles, is home to 1.6 million people. To do so I introduce a new type of data – real-time, geocoded measures of micro-level urban foot traffic extracted from camera feeds – that is becoming increasingly available for social science researchers. The data, which I refer to as "foot traffic data" throughout the remainder of the paper, will soon be available for large swathes of Manhattan which, as noted above, is an optimal laboratory for new research on civic and political behavior.

The data offer numerous new opportunities for understanding *context* and how it shapes opportunities for civic and political participation. It can provide insights into everyday political and social behaviors, but can also furnish a novel perspective on large public events, such as rallies and protests. Importantly, the data can be combined with other sources and types of data, for example pedestrian and public transportation infrastructure, citizen relationship management data (such as NYC's 311 service), polling locations, crime reports, and so forth. In doing so, researchers can gain better cognizance of how behavior and space intersect. Finally, foot traffic data provides new methodological challenges for social scientists working with micro-level spatial data, some of which I address here.

The foot traffic data utilized in this paper are obtained from a company called Placemeter. Placemeter measures activity on streets and sidewalks from a network of sensors, comprising mostly publicly and privately owned cameras.[2] The startup uses computer vision technology to convert video feeds into counts of pedestrians, cars, and cyclists, generating a continuous feed of information about street-level behavior. The company's data is being used by businesses to, for example, make decisions about commercial real estate (where to locate a store or place a billboard), or to inform potential customers about expected wait times at a store. It is also being utilized by local government agencies to better understand pedestrian traffic flow in parks and public spaces, improving public policy in cities. For instance, the Greenpoint, Brooklyn Chamber of Commerce is using Placemeter's data to determine the optimal location for trash cans in the neighborhood.

The remainder of the paper is organized as follows. In section 2.2 I provide a detailed introduction to the pedestrian traffic data. In particular, I describe how the data is generated, its structure, and a number of features that make it unique in social science settings. Next I put forth empirical applications relating to different types of behavior: mass public protests, 311 reporting, and crime. In section 2.3 I show how this new data can be used to study changes in foot traffic patterns throughout an urban setting when mass political protests and mobilizations occur. This provides a unique snapshot into the behavior of citizens during mass political events. In section 2.4 I map foot traffic to two other features of city life. First I demonstrate how the data can be used to study the relationship between 311 reporting, which is increasingly being used to proxy for various civic and political behaviors (Lerman and Weaver 2014; Levine and Gershenson 2014; O'Brien et al. 2016), and citizen flows. Next I show that more foot traffic tends to mean more major crime incidents. Then, in section 2.5, I provide a synopsis of future applications of the data, some potential and some already underway. Section 2.6 concludes.

[2] The network comprises existing public camera feeds, including NYC Department of Transportation (DOT) traffic cameras, Placemeter-owned IP cameras, and Placemeter sensors. Placemeter sensors are primarily old, unwanted smartphones that are being put to use to as cameras. These re-used smartphones belong to city residents who convert their phones into street sensors for Placemeter in exchange for a monthly cash payment. The company provides participanting residents a kit, including a window suction cup, to convert their unused smartphone into a data-capturing device. The better the view, the more the participant gets paid– up to $50 per month.

2.2 Data

The raw foot traffic data are a spatial time series: a set of geo-referenced points with corresponding dates, times, and counts. The counts measure foot traffic – how many people are in any given space at any given time. The foot traffic counts are based on visual data that is parsed by proprietary algorithms that take into account the number of people observed per second, estimated walking speed, and the average length of "dwell time" (how long an individual remains in the frame). The methodology generates highly accurate estimates of the level of foot traffic for the given time interval. Because camera feeds are continuous, the data are available at various time intervals, most often 15-minute counts or 1-hour counts, which can be aggregated to produce average daily foot traffic figures. Further, the counts are either directional foot traffic counts (see application 1) or non-directional foot traffic counts (as used in application 2).

The first set of data, from the spring of 2015, utilizes a "virtual turnstile" method to detect the direction of pedestrian walking patterns at 15-minute intervals. To understand the intuition behind this particular methodology, imagine a virtual turnstile in the field of view of each sensor that counts people crossing it in both directions. Each sensor detects movement in two directions, depending on the orientation of the street and the camera. For example, a sensor placed on Fifth Avenue in Manhattan, which runs roughly north to south, counts the number of pedestrians moving north (uptown) and the number of pedestrians moving south (downtown) at any given point in time.[3] Meanwhile, a sensor located on 34th Street, which runs roughly east to west in Manhattan, is able to detect pedestrian movement east (towards the East River) and west (towards the Hudson River). The second set of data utilized in this paper contains non-directional hourly foot traffic counts from the summer of 2014, prior to the roll out of the virtual turnstile technology.

[3]Technically, the avenues of Manhattan do not run perfectly north-south because the grid plan is aligned with the Hudson River. Many tourist maps, as well as the NYC subway map, rotate the city such that the grid follows the cardinal directions. In this paper, I refer to the cardinal directions for the sake of simplicity.

Data Limitations

Although the data is being continuously improved, there exist two main limitations: sensor error and camera placement. Below I discuss the consequences of each, and how I attempt to mitigate them, in turn.

Sensor error can arise in a few different ways. Though camera locations are generally fixed, they can occasionally move. When this occurs, a sensor will give erratic counts; for example, if the virtual turnstile, or measurement point, moves to focus on, say, a building wall rather than a street, the sensor will record (inaccurately) a string of zero counts for foot traffic. A second, but less common, form of sensor error is connectivity interruptions. Sensors, because of a dropped or faulty Internet connection, can go down from time to time, producing zero counts at certain hours. Both of these problems are being addressed in future generations of these data; for instance, new sensors will use cellular networks to transmit data. In the analyses presented throughout this paper, I use only the most reliable sensors, as indicated by Placemeter staff. Further, sensors that suddenly record zero traffic for an unrealistically long period of time are dropped from the data for the entire day on which the inaccurate recording took place.

The second major limitation of the foot traffic data is that coverage is limited by the physical placement of cameras. As noted above, Placemeter utilizes its own IP cameras and sensors to supplement public camera feeds, but the placement of those private devices is in large part determined by the needs of their clients, who are typically interested in just a few discrete measurement points, rather than an entire neighborhood. Since these points are not evenly distributed throughout the city, in my cases I restrict my analysis to neighborhoods, such as Midtown, where coverage is more comprehensive. Again, as the data become more reliable in the coming months and years, the number of camera feeds will increase, improving the coverage throughout urban spaces.

To make the most of the available data, it is key to move from the network of discrete camera placement points to a continuous surface that represents traffic flow. To do this I use spatial kriging, a method of interpolation that estimates a value for a given unmeasured point based on a weighted average of the values of surrounding measured points, called *control points*. Kriging estimates the optimal set of weights for each unsampled location based on the relationship between each of those points and its surrounding control points. Weights are assigned according to a data-driven weighting function that takes into account spatial covariance values. Based on basic linear regression, kriging minimizes error variance. Thus, one advantage of kriging over simpler interpolation methods, such as inverse distance weighting (IDW), is that it is possible to quantify the error associated with each estimated value(O'Sullivan and Unwin 2014). After plotting the estimation variance, I restricted the main analysis to neighborhoods in which there is a relatively high degree of certainty in the interpolation.[4]

One final limitation worth noting is that the earlier (2014) data used in the second application are unreliable after dark. Thus, the figures reported in section 2.4 are for the hours between morning nautical twilight and evening nautical twilight (roughly 6:00am to 8:00pm) based on historical data on sunrise and sunset times by day. Placemeter's methodology for processing data before dawn and after dark has subsequently improved, and thus the spring 2015 dataset includes all hours of the day.

In summary, after applying spatial kriging to the data, the result is a time-stamped smooth surface of traffic flow counts that cover much of Manhattan during, in 2014, daylight hours, and in 2015, all hours of the day. This data layer can then be easily merged to other geo-referenced (and time-stamped) data. In the next two sections, I demonstrate how these data can be used to approach classical questions in political and social science.

[4]Kriging was done using the `geoR` package in `R` (Ribeiro Jr. and Diggle 2001). In order to obtain the parameter inputs for the kriging procedure, the spatial variation was first described using a semivariogram. This plot appeared to be best fit by an exponential function, indicating that the distance decay in dependence should be represented by an exponential model in the kriging procedure.

2.3 Application 1: Protests and Pedestrians

The everyday pulse of cities is sometimes punctuated by extraordinary episodes like protests and rallies. The dynamics of mass protests have primarily been studied via surveys and interviews (e.g. Dionne et al. 2014), though attempts to quantify crowds with mobile phone and twitter data are increasingly common (Botta, Moat, and Preis 2015). Foot traffic data provide a unique tool with which to study the way protest actions and rallies unfold in urban settings. For example, the data can be used to understand how crowds develop, whether protesters arrive from one direction or multiple directions, or the circumstances under which crowds appear all at once, versus gradually over time.

I use directional counts recorded by a set of well-placed sensors to study how pedestrian traffic changes during a major protest event. Directional data allows for a nuanced look at pedestrian flows; each sensor records foot traffic counts in two cardinal directions. Some cameras are positioned so that they record north-south movement, while others record only east-west movement. To account for this, the kriging procedure (which maps the point data to a continuous surface) is conducted separately for both types of directional foot traffic.

To better understand the directional data, consider figures 2.1 and 2.2 below. These figures show directional (south to north – toward uptown) foot traffic patterns in 15-minute intervals from 5:00pm to 7:00pm, and from 7:00am to 10:00am, respectively, on a selection of weekdays in the Spring of 2015. Comparing the two figures, the most immediate observation is that foot traffic towards uptown is heavier during the evening rush hour than it is in the morning rush hour. Sensors in Midtown detect relatively high numbers of pedestrians during both rush hour periods, but there is also a spike in traffic around Union Square and Greenwich Village that occurs in the evening but not in the morning. This provides face validity because these neighborhoods tend to be after-work destinations, where residents and visitors go to shop, eat, and drink. Note that the median foot traffic in one direction does not exceed 90 people per 15 minutes below Midtown, even in heavily-trafficked neighborhoods like those around Union Square and 14th Street.

Figure 2.1: Evening rush hour: Interpolated foot traffic in the direction of uptown in 15-minute intervals from 5:00pm to 8:00pm on weekdays (Monday, Wednesday, and Friday). “S” indicates location of functional sensor; sixteen sensors that record south-to-north foot traffic were functional during these times. Dates include April 7th, 14th, 21st, 22nd, 24th, and 29th, and May 1st, 6th, and 8th, 2015.

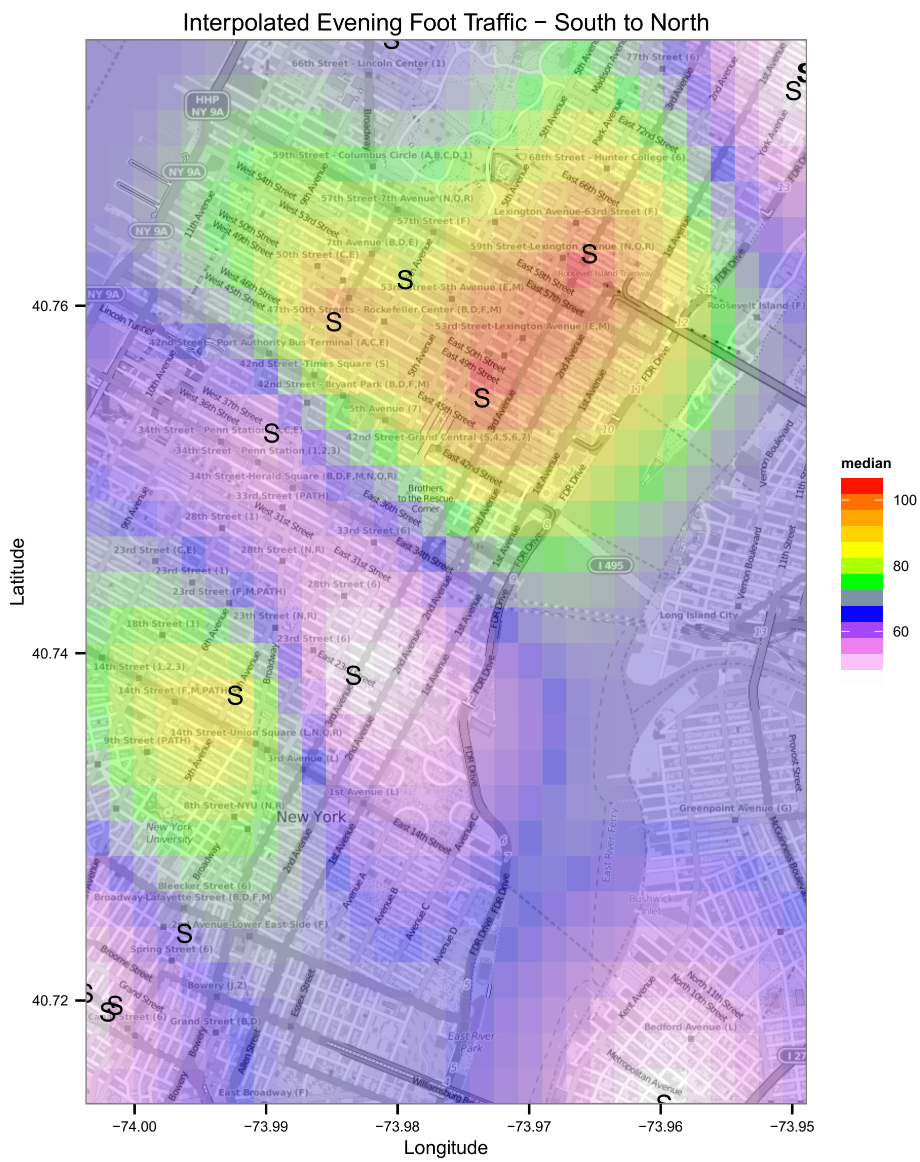

Figure 2.2: Morning rush hour: Interpolated foot traffic in the direction of uptown in 15-minute intervals from 6:00am to 9:00am on weekdays (Monday, Wednesday, and Friday). "S" indicates location of functional sensor; sixteen sensors that record south-to-north foot traffic were functional during these times. Dates include April 7th, 14th, 21st, 22nd, 24th, and 29th, and May 1st 6th, and 8th, 2015.

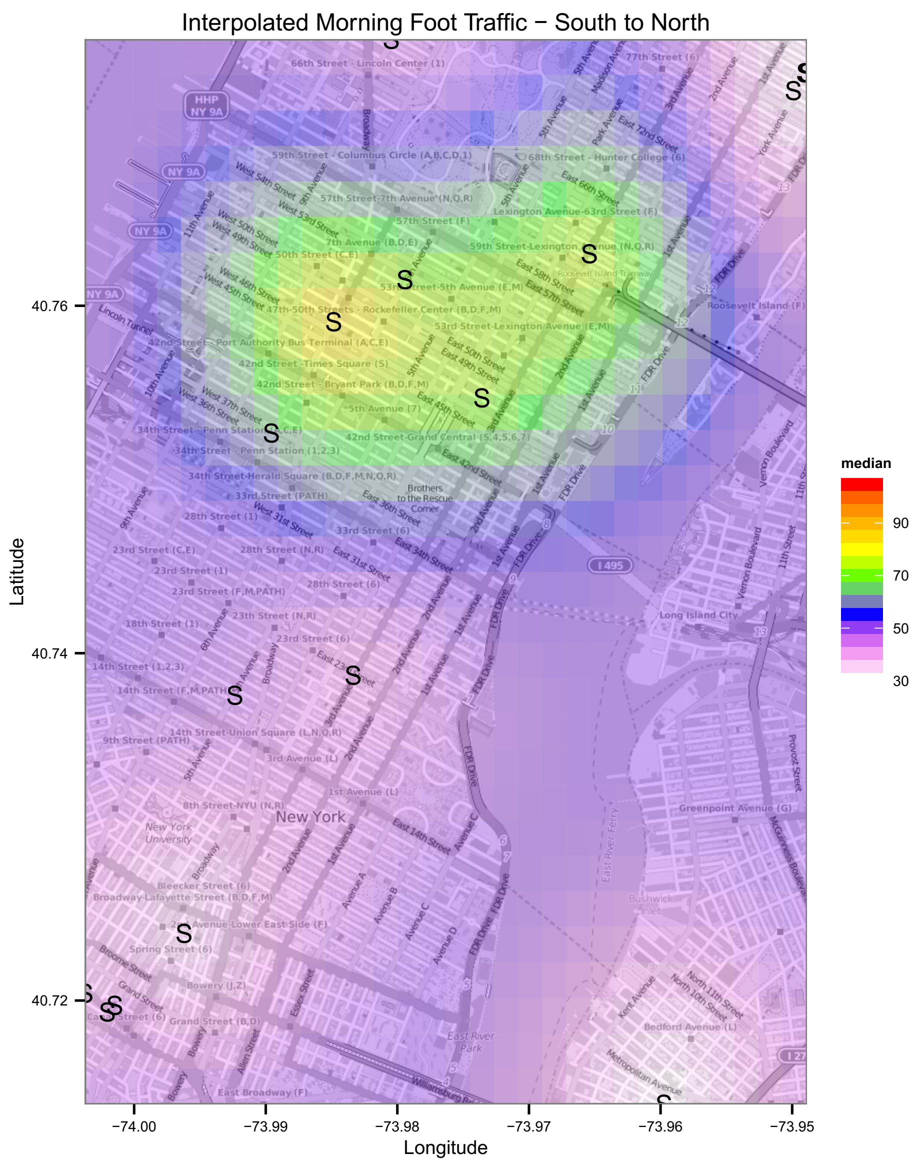

In April and May 2015, in response to the death of Freddie Gray while in police custody in Baltimore, Maryland, rallies and protests were held in cities around the United States. At the time, this dramatic collective social action was the culmination of a number of cases of perceived police violence against African Americans, starting with the killing of Michael Brown in Ferguson, Missouri, in 2014. On April 29th 2015, a major rally was held in New York City. Publicized widely via social media, the event was scheduled for 6:00pm in Union Square. I find clear surges of traffic towards the Union Square park, as well as away from it, that correspond with the ebb and flow of a protest.

The data I consider were recorded by a camera positioned on the northern sidewalk on 14th Street between Union Square West and Fifth Avenue, as shown in figure 2.3. This location, near the southwest corner of Union Square Park, is on a major thoroughfare leading towards and away from the park. The park also happens to sit on top of an important intra-city transit hub and transfer point; the 14th Street - Union Square subway station is the fourth-busiest station in NYC (after stations located at Times Square - 42nd Street, Grand Central - 42nd Street, and 34th Street - Herald Square), with an annual ridership of 35,677,468 in 2014.[5] Thus the foot traffic counts reflect daily commutes, hence the spikes during mornings and evenings, as well as a substantial amount of "off-peak" traffic.

In figure 2.4 I consider foot traffic data measured at 15 minute intervals, on April 29th 2015, the day of the protest. I then systematically compare the foot traffic patterns to a "placebo" day, April 22nd 2015, one week before the protest. Deviations from the foot traffic patterns observed on the 22nd can be interpreted as symptomatic of what was unusual during the protest. That is, we can compare a "normal" Manhattan Wednesday to the exceptional Wednesday. I find that the densities of foot traffic for the protest day and the placebo day are similar in the morning, mid-day, and late at night, but diverge strikingly in the afternoon and evening.[6] These divergences are explored in more detail below.

[5] "Facts and Figures: Annual Subway Ridership". Metropolitan Transportation Authority. Retrieved 2015-08-07. `http://web.mta.info/nyct/facts/ridership/ridership_sub_annual.htm`

[6] Warmer weather on April 29th (a mean of 63 degrees Fahrenheit and a high of 75 degrees, versus a mean of 56 degrees and high of 63 degrees on the previous Wednesday) could be driving the overall elevated levels

Figure 2.3: Map of vicinity of Union Square. The red square indicates sensor location, on 14th Street between 5th Avenue and Union Square West.

Figure 2.4: Foot traffic in (eastbound) and out (westbound) of Union Square on Wednesday, April 29th, 2015 vs. Wednesday, April 22nd, 2015. Sensor location is on 14th Street between 5th Avenue and Union Square West.

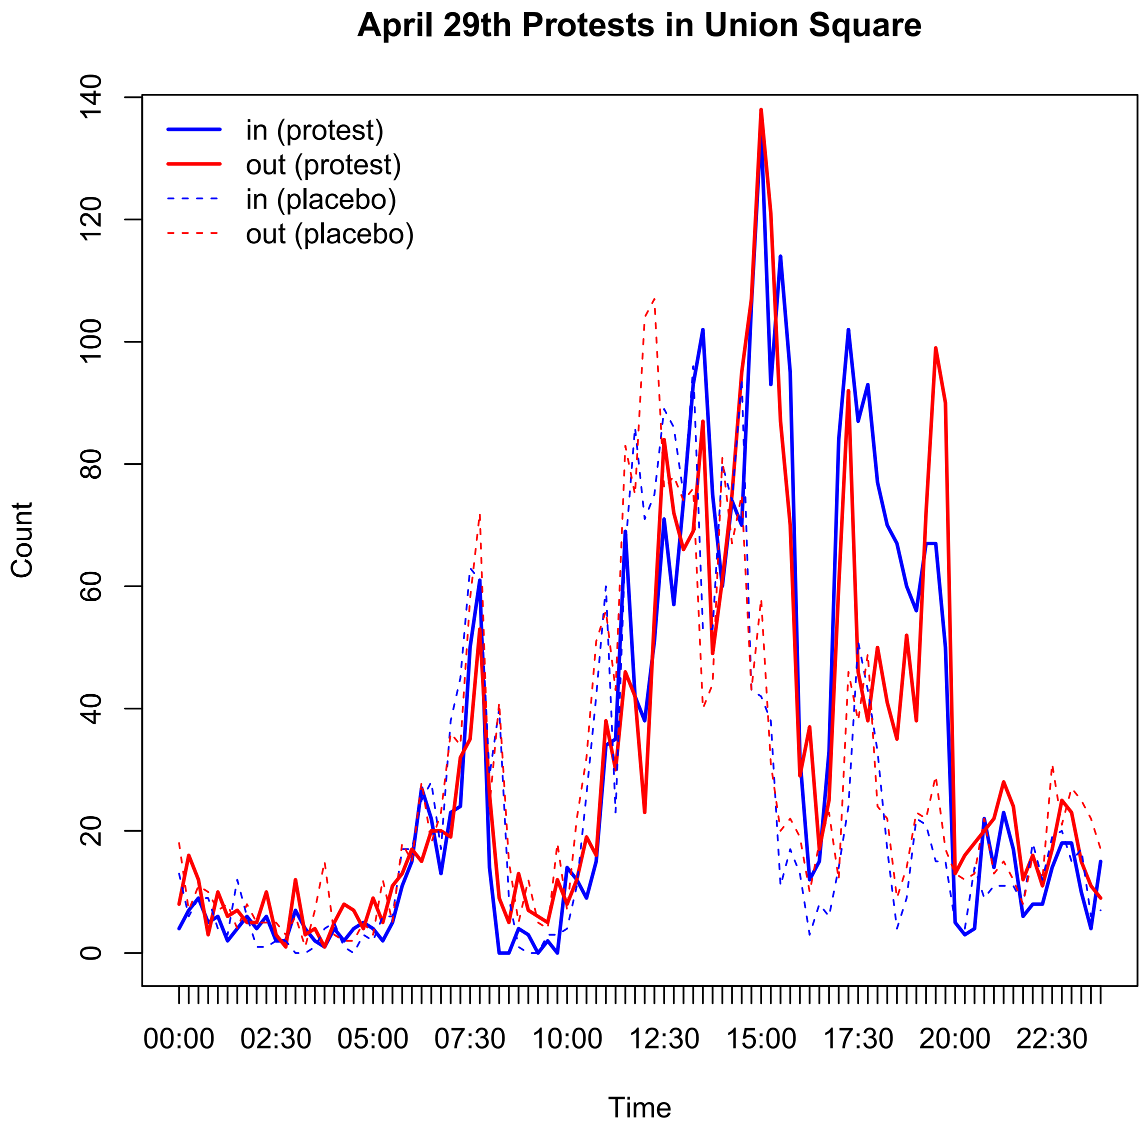

As seen in figure 2.4, foot traffic is extremely similar to the placebo day during the morning, the early mid-day, and late night. However, foot traffic in both directions is unusually high around 5:15pm, with higher levels of people streaming in than out until around 7:00pm, when the traffic flow switches direction. Although a portion of this traffic is due to individuals commuting to and from work at the end of the day, the spike on April 29th is much higher than it is during a typical Wednesday rush hour, as can be seen in the distance between the solid and dashed lines, which represent the placebo day. This is consistent with a rally scheduled at 6:00pm that concluded at 7:00pm, when protesters spread out across the city.[7] Estimates of rally attendance ranged from "hundreds"[8] to "thousands"[9] of demonstrators. As Union Square Park has multiple pedestrian entry and exit points, the Placemeter sensors on 14th Street would have captured only a fraction of the total foot traffic to and from the rally. As such, the measures should be seen as an assessment of the ebb and flow of the protest, not measures of the actual magnitude of the event. If more cameras were available, it may be possible to back out accurate estimates of attendance numbers for major events.

Of course, the analysis of rally-related foot traffic is inherently constrained by the placement of cameras that record pedestrian traffic.[10] Since the 2015 data include just one sensor in the immediate vicinity of Union Square, one can only make reliable inferences about foot traffic on a limited portion of the sidewalk. In the near future, however, I hope to have access to a more comprehensive network of sensors.

---

of foot traffic in the afternoon. However, the warmer weather does not explain the divergent trends in and out of Union Square on the evening of the rally. Precipitation did not occur on either day.

[7] `http://newyork.cbslocal.com/2015/04/29/mayor-de-blasio-nyc-protest-baltimore/`

[8] `http://www.crainsnewyork.com/article/20150429/POLITICS/150429811/hundreds-rally-in-union-square-protesting-death-of-freddie-gray`

[9] `http://www.myfoxny.com/story/28926801/solidarity-protest-for-baltimore-to-take-place-in-nyc`

[10] It is also constrained by the dates for which I was able to obtain data. At this point in time, Placemeter has only provided me a limited subset of data.

2.4 Application 2: Sentinels of the Sidewalk

The sidewalks of New York City were a favorite subject of writer and activist Jane Jacobs. In *The Death and Life of Great American Cities* (1961), Jacobs argued that "eyes upon the street" – by which she meant the watchful gaze of neighborhood shop-owners and residents, alert to minor disturbances and problems – are essential to a vibrant, safe city.

Broken Windows

In 2003, the New York City's "eyes" gained a direct link to city government with the launch of NYC 311, NYC's non-emergency customer service line. NYC's 311 line operates 24/7 and assists 60,000 callers per day.[11] Residents and visitors report everything from noise complaints to broken parking meters to rodent sightings. Those who contact 311 can be thought of as "sentinels of the sidewalk", supplying vital information about what is happening on the ground to the city agencies responsible for fixing problems.

But was Jacobs right? Does higher foot traffic, and therefore more "eyes on the street", translate into more 311 reports, or more active citizen engagement? A high density of pedestrians may mean greater potential for complaints about a blocked sidewalk, an overgrown tree branch, or an overflowing litter basket. Yet perhaps there are diminishing returns to eyes on the street, and perhaps crowds even hinder the ability of sidewalk sentinels to do their work. Using data from Placemeter and NYC OpenData,[12] one can study how rates of pedestrian traffic dictate what gets reported to 311 by citizens.

First, I combined NYC 311 reports with the foot traffic data from over the course of two weeks in the summer of 2014. On top of the foot traffic layer, generated via spatial kriging as outlined in Section 2.2, I then overlaid NYC 311 data from NYC OpenData for the same period and extracted predicted values of hourly foot traffic for each geocoded 311 report. In

[11] `http://www1.nyc.gov/311/index.page`

[12] `https://nycopendata.socrata.com/Social-Services/311-Service-Requests-from-2010-to-Present/erm2-nwe9`

other words, I construct a dataset of 311 calls, where each 311 call is linked with how many pedestrians per hour, on average, were in the vicinity of the problem when it was reported to the City.

Figure 2.5 visualizes the data as described. Each point corresponds to a 311 report related to outdoor spaces – issues with sidewalks, streets, parks, or playgrounds. Each point is colored according to the average amount of foot traffic in the vicinity, as measured by the hourly foot traffic data after kriging. Readers familiar with New York City will not be surprised to see that the highest levels of foot traffic are to be found between 34th Street/Penn Station and 42nd Street/Times Square between Eighth and Fifth Avenues, traditionally high volume areas. Foot traffic diminishes as one moves further towards the Hudson River, hitting a low point of several hundred people per hour around the more industrial Hudson Yards. Pedestrian rates are also lower on the East Side, from Midtown East through Murray Hill and Kips Bay, areas with less tourist appeal than the bustling center of the city.

In terms of inferring patterns in the data, a more useful visualization is presented in figure 2.6. This figure provides density curves showing the relationship between foot traffic and various types of 311 complaints in Midtown for June 17-30, 2014. The average distribution of hourly foot traffic across Midtown is plotted in gray, and overlaid with density curves corresponding to several categories of 311 reports from the same time period. If a call type occurs uniformly across the city (i.e., if there is no relationship between foot traffic and call frequency) then the colored line would track the dotted gray horizontal line at zero. Thus the distance between a colored line and the dotted line at any given point represents the deviation, negative or positive, between the number of calls that we would *expect* to observe under the null hypothesis of no relationship and the number of calls of that type that we actually observe. So, for example, noise complaints (in red) occur less than expected in high-traffic areas and more than expected in middle to low-traffic areas, as would be anticipated.

The first take-away from the visualization in figure 2.6 is that, indeed, more people means more "eyes on the street", but only up to a point: when foot traffic is at its highest,

Figure 2.5: On this map of Midtown Manhattan, each dot corresponds to a 311 report pertaining to an issue occurring outdoors– on a sidewalk, street, park or playground– during last two weeks of June, 2014. The colors of the dots correspond to the average amount of foot traffic, measured in people per hour, at the location of the 311 report.

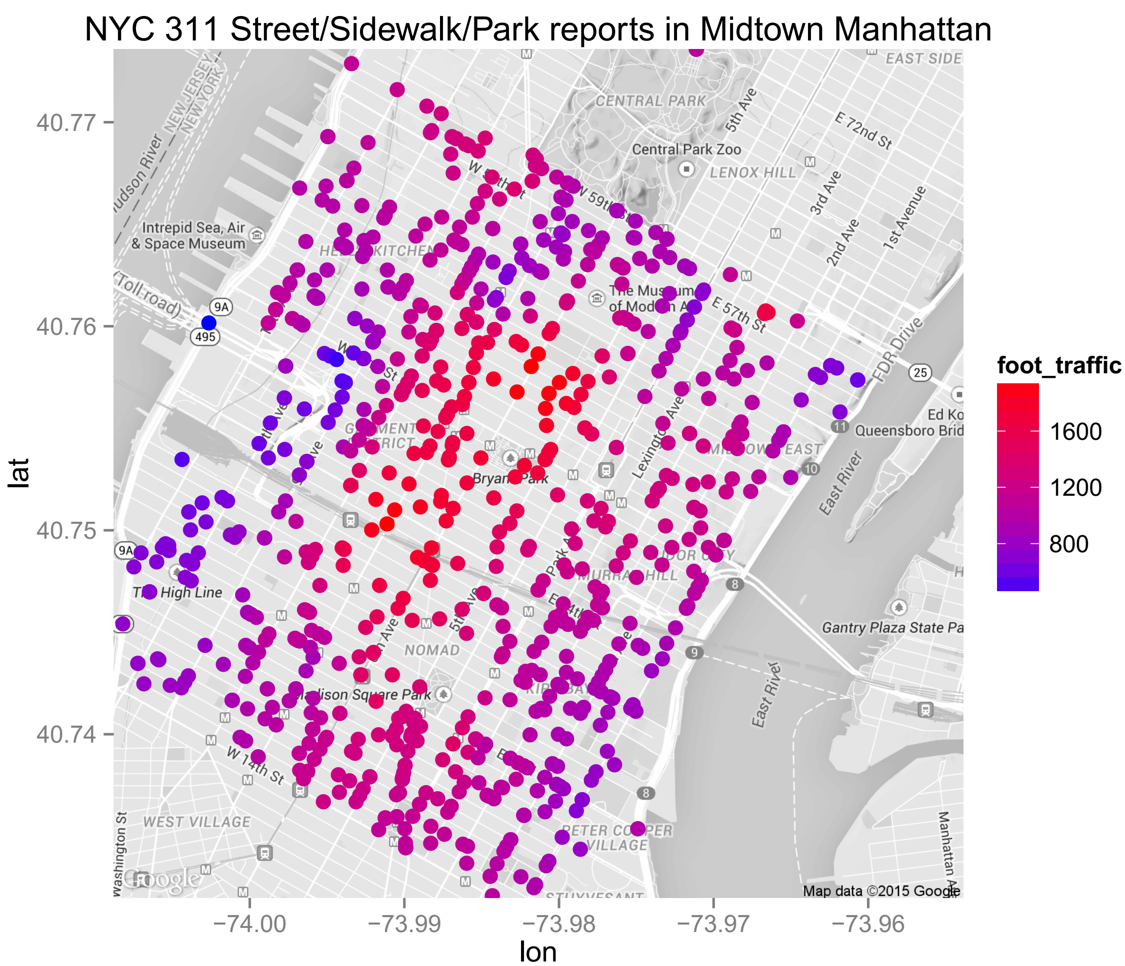

Figure 2.6: This density plot shows the relationship between hourly foot traffic and various types of 311 complaints in Midtown for June 17-30, 2014. The average distribution of hourly foot traffic across Midtown is plotted in gray, and overlaid with density curves corresponding to several categories of 311 reports from the same time period.

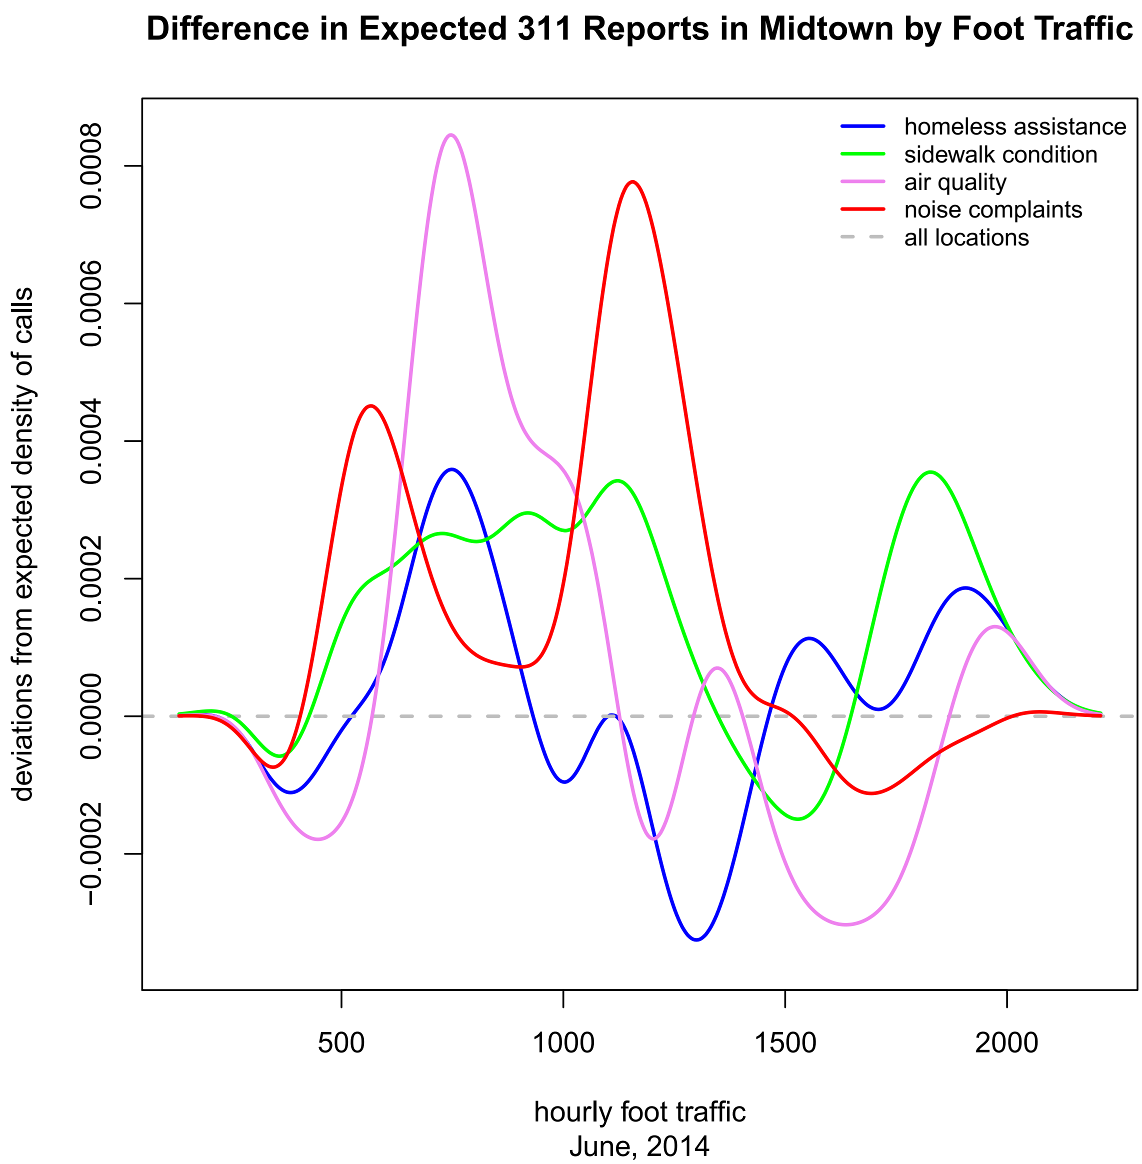

311 reporting tends to begin to drop off. One possible explanation is that the type of people on the street matters; highly trafficked areas are often those that swarm with tourists, who would obviously be less likely than residents to report problems to 311. Relatedly, stretches of pavement near major intra-city public transit hubs, such as the 42nd Street subway station, seethe with pedestrians on their way to and from various destinations, who are unlikely to stop and consider their surroundings long enough to file a 311 complaint. Another potential explanation is the bystander effect, a social psychological phenomenon in which the presence of many other individuals decreases the chances that any one person will report a problem (Darley and Latane 1968). In all likelihood, some combination of these processes are at work.

The second take-away is that the relationship between foot traffic and 311 reporting varies by the type of complaint filed, as depicted by the different colored lines in figure 2.6. To determine whether the distribution of each 311 call type by hourly foot traffic is indeed statistically distinguishable from the null distribution, i.e., the distribution of calls we would have observed if call counts were proportional to foot traffic, I conducted a series of two-sided Kolmogorov-Smirnov tests (Wang, Tsang, and Marsaglia 2003). These tests indicate that 311 reports pertaining to both homelessness and noise complaints differ significantly from the baseline distribution ($p < 0.05$), while the distributions of sidewalk condition ($p \approx 0.73$) and air quality complaints ($p \approx 0.21$) do not. This may suggest that fundamentally different processes underscore different types of reporting.

These results imply that for certain categories of 311 complaints, the number of calls is relatively unaffected by the quantity of foot traffic; complaints about sidewalk condition and air quality are generated at similar rates regardless of how many people are present. Meanwhile, noise complaints and calls for assistance for the homeless follow different patterns. On average, street noise is rarely reported at the most bustling Midtown locations. This may reflect expectations about where noise is acceptable, e.g., we expect Times Square to be loud but are less tolerant of noise on a residential block. Meanwhile, reports of homeless persons in need of assistance are rarely made in the least trafficked areas but are over-represented in places with a high density of foot traffic.

An important limitation is that we do not have a record of *who* is filing these reports. A recent study of 311 users in Boston, Massachusetts finds that reporting issues in the public domain is associated with territorial motives, or the desire to claim ownership and responsibility for spaces (O'Brien et al. 2016). If a similar dynamic is at work in Midtown Manhattan, it is likely that the bulk of 311 reports are generated by residents and, perhaps, local business owners, rather than by tourists or by so-called "bridge and tunnel" commuters. Furthermore, it seems very plausible that complaints about noise or air quality would likely come from individuals who live in a neighborhood, rather than those simply passing through. Thus 311 reporting behavior as a function of foot traffic can perhaps be understood as a way in which New Yorkers exercise some degree of control over sidewalks that they share with a much wider populace; at the same time, their beliefs about what is a tolerable degree of disorder – as well as what types of disorder are acceptable and where – adjust depending on the neighborhood or setting.

Major Crime

Though the relationship between foot traffic and minor disturbances is mixed, the link between the former and serious crime is more stark. The New York City Police Department recently made public quarterly reports of seven major felonies at the incident level. These include robbery (theft with force or threat of force), burglary (breaking and entering, entering into any structure with the intent to commit a crime inside), grand larceny (any theft of property where the value of the property is greater than $1,000), felony assault (assault that causes physical injury). Using these data, I mapped incidents of crime occurring in the first half of 2014, and in the two week period for which foot traffic data are available.[13]

[13] According to the NYPD, "Crimes occurring anywhere other than at an intersection are represented by a midpoint X coordinate and a midpoint Y coordinate (center of block). Rape offenses are geocoded as occurring at the police station house within the precinct of occurrence. Offenses that lack an X coordinate and Y coordinate are geocoded as occurring at the police station house within the precinct of occurrence. Offenses occurring in open areas such as parks or beaches may be geocoded as occurring on streets or intersections bordering the area."

Visual inspection and hotspot analysis reveal a striking pattern: major crime incidents in Manhattan cluster in midtown. Figure 2.7, which maps the occurrence of the seven types of major felonies in all of Manhattan for the period June 17-30, 2014, shows that the parts of the island most frequented by tourists are also the most dangerous. The most common crime is grand larceny, but a similar pattern emerges even when dropping that incident type from the analysis.[14]

Due to foot traffic data availability, the following analyses focus on the middle section of Manhattan, above 14th Street and below 59th Street. This area not only has the most comprehensive traffic camera coverage, and thus the most reliable interpolated values; it also happens to contain regions of some of the highest crime incidence in the city. Though I will refer to this area as Midtown for short, it encompasses the neighborhoods of Clinton, Hell's Kitchen, Chelsea, the Garment and Theater districts, Murray Hill, Kips Bay, and Gramercy Park. This five-square-mile area is home to a total of 310,774 people, around one fifth of Manhattan's population.[15]. This figure is roughly equivalent to the number of pedestrians who enter the heart of Times Square every day.[16]

Major crime is positively correlated with foot traffic, both across all of Manhattan and within midtown. The left panel in figure 2.8 shows each type of offense plotted over interpolated foot traffic values. The darker red shading indicates higher levels of pedestrian traffic. Crime incidents appear to cluster towards the center of the map, in areas of darker shading. The right panel uses the same data to display the contours (in blue) of a crime hotspot map, produced via two-dimensional kernel density estimation. Analogous to a topographical map, contour lines join points of equal crime frequency, as if they are land features of equal elevation. The shading of the lines, from dark to light blue, indicate levels of crime

[14]NYC does not make available geocoded incidents of lesser crimes, only the seven major felonies.

[15]This area contains 205,134 housing units – almost a quarter of those in all of Manhattan. In 2010 the borough was home to 1,585,873 people, in 847,090 housing units. Source: U.S. Census 2010.

[16]An estimated 300,000 people enter each day, on average, according to the Times Square District Management Association: http://www.timessquarenyc.org/do-business-here/market-facts/pedestrian-counts/index.aspx#.V35gxesrJFE

Figure 2.7: Crime types in Manhattan

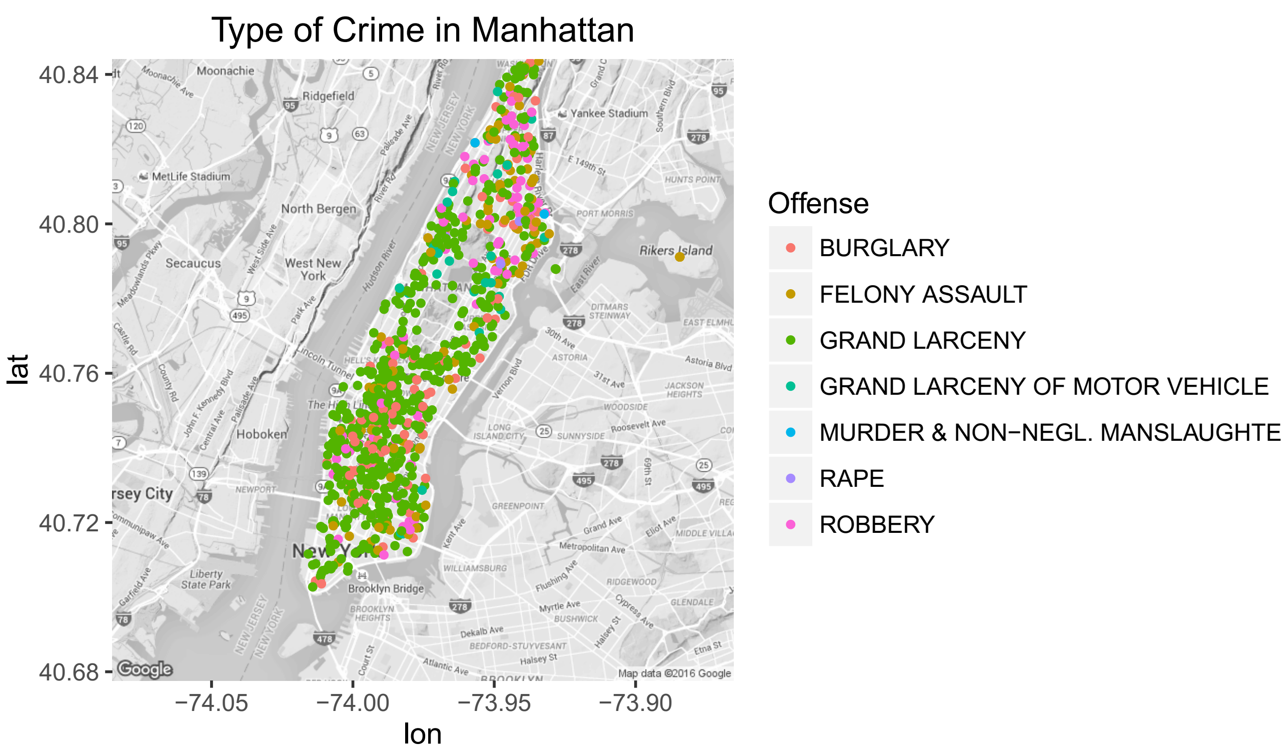

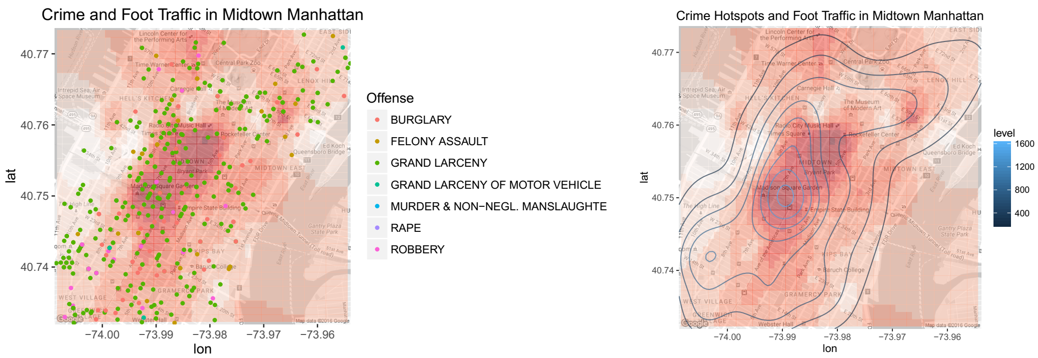

Figure 2.8: Crime and foot traffic in Midtown

from low to high. Peak crime intensity appears to be around 34th Street between Sixth and Seventh Avenues, a few blocks below the heart of Times Square.

As further evidence of a positive association between foot traffic and crime, consider figure 2.9. The interpretation is analogous that of figure 2.6 above: if there is no relationship between foot traffic and crime then the blue line would lie on top of the dotted gray zero line. The overall pattern suggests pedestrian density and crime density move roughly in tandem.

In order to account for population characteristics which may confound the relationship between crime and foot traffic, I control for various neighborhood features in a Census block level regression analysis. To do so I aggregate crime incidents up to the Census Block, each of which corresponds roughly to a Manhattan city block, and create an analogous measure of foot traffic by calculating the mean interpolated foot traffic for each block. I merge this block level dataset (n = 852) with population counts by race and ethnicity from the U.S. Census. Census blocks are the smallest unit at which population data are available; while aggregating entails the loss of some detail, with 365 people on average per block these units are still quite small.

In table 2.1 I regress the count of crime incidents per block on mean block-level foot traffic. For ease of interpretation, OLS regression coefficients are reported; the results are

Figure 2.9: This density plot shows the relationship between hourly foot traffic and major crime incidents in Midtown for June 17-30, 2014. The average distribution of hourly foot traffic across Midtown is plotted in gray, and overlaid with a density curve corresponding to major crime incidents from the same time period.

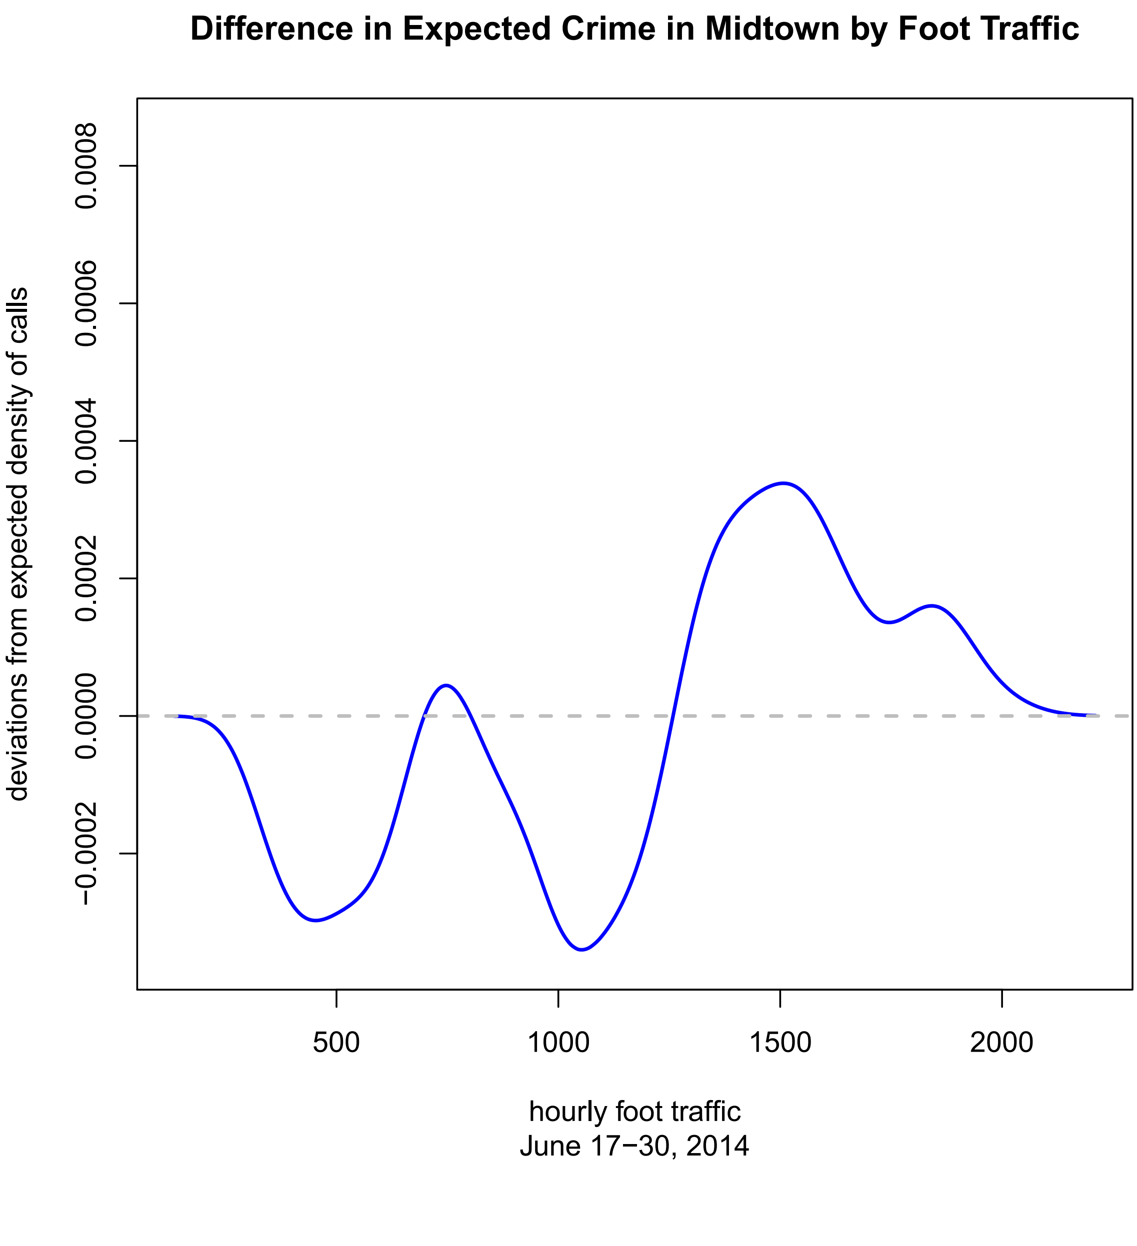

substantively similar when Poisson regression is used to account for the count nature of the dependent variable. Columns 1 and 2 include controls for the racial and ethnic make-up of the block, population, and number of housing units. Column 3 includes Census tract fixed effects, which hold constant fixed characteristics of each of the 77 tracts represented in the sample. The latter column, in which coefficient estimates pick up variation that comes only from *within* each of those tracts, represents a very conservative comparison between blocks; Census tracts are relatively homogeneous units with respect to population characteristics and economic status. These tracts contain nine Census blocks on average, and correspond to an area that residents are likely to walk through on a daily basis, on their way to and from public transit hubs. Even within these small areas, the positive relationship between crime counts and foot traffic remains unchanged. Based on column 3, on average an increase of 100 pedestrians per hour is associated with 0.1 increase in the number of serious crime incidents.

Table 2.1: Major crime and mean foot traffic

	Crime count		
	(1)	(2)	(3)
Foot traffic	0.001^{***}	0.001^{***}	0.001^{***}
	(0.0001)	(0.0001)	(0.0003)
Constant	-0.264^{**}	-0.397^{***}	-0.692^{*}
	(0.130)	(0.139)	(0.370)
Controls?	No	Yes	Yes
Fixed effects?	No	No	Yes
Observations	697	697	697
Adjusted R^2	0.045	0.065	0.155
Note:			*p<0.1; **p<0.05; ***p<0.01

Table B.1 in the Appendix repeats this analysis replacing the dependent variable with crime incidents per capita. The 274 Census Blocks with zero population are dropped, leaving 512 blocks in total. The results are substantively similar in the first two columns: an increase of 100 pedestrians per hour is associated with a 1 percentage point increase in crime per capita, on average. When accounting for Census tract fixed effects, however, the relationship is indistinguishable from zero. This exception aside, overall the correlation between foot traffic and felonious crime tends to be positive and statistically significant. This is true even accounting for population characteristics and/or dropping non-residential Census blocks from the analysis.

Discussion

Nearly a half century after Jacobs coined the term "eyes on the street", the combination of foot traffic data with 311 and crime data allows us to explore the nuances of an idea that has influenced urban planners around the world. These novel data sources can breathe new life into Jacobs' ideas about the vitality and safety of the city and its sidewalks. They allow us to study the everyday relationship between citizens, cities, and political and civic behaviors.

Here I find mixed support for Jacobs' claims. Eyes on the street correlate with 311 reporting, but they can be overwhelmed by excessive foot traffic. Serious crime, meanwhile, is negatively associated with the number of eyes on the street. This finding, on the face of it, appears to run counter to Jacobs' intuition about neighborhoods with high levels of pedestrian traffic. However, key to Jacobs' assertion of the protective benefits of "eyes on the street" are the people who live and work in a neighborhood and are committed to its well-being. Perhaps residents of Midtown, whose blocks are besieged day and night by tourists and visitors, are unable to effectively monitor their turf. This may occur because residents intentionally block out their surroundings to cope with over-stimulation (Simmel 1978), or because there is simply too much going on to notice something abnormal or suspicious.

Examining figure 2.6 alongside figure 2.9 provides suggestive evidence that high levels of foot traffic, particularly around 1,500 people per hour, are detrimental to any protective benefits of eyes on the street. While major crime spikes, 311 calls tend to drop off, relative to levels in less-trafficked locations. In Midtown, there are no deserted streets like the ones

Jacobs warned us about. Rather there are, perhaps, streets so crowded that safety and civicness give way to crime and disorder.

Though the evidence presented here is speculative, it suggests avenues for future research. The degree to which urban denizens function as 'eyes' is context-dependent. New data sources – like those utilized here – create opportunities for a more nuanced engagement with street-level phenomena.

2.5 Next Steps and Future Applications

Modeling pedestrian movement has historically been the purview of planners designing transportation systems, shopping centers, city sidewalks, and evacuation routes, with an emphasis on keeping vehicular and pedestrian traffic separate (Fruin 1971; Pushkarev 1975). For instance, the timings of traffic control devices, such as pedestrian "Walk" or "Don't Walk" signals, and the locations of crosswalks and elevated walkways, are often informed by manual foot traffic counts (Emmons 1965; Løvås 1994). In addition, commercial enterprises have leveraged such information in retail settings in order to convert traffic into sales (Perdikaki, Kesavan, and Swaminathan 2012). Yet the myriad ways in which understanding foot traffic patterns can augment social science inquiry – like the two applications introduced in this paper – have been largely ignored. A number of examples of other studies, some potential and some in preparation, are described below.

First, foot traffic data has the potential to address a host of unanswered questions about how cities shape political and civic behavior. For example, how does foot traffic change on Election Day, when pedestrian patterns are shaped by where and when residents vote? Voting, which is the quintessential political act, may heighten citizens' civic engagement with local government. In one future project, I will investigate how 311 reporting changes on and after Election Day, conditional on shifts in pedestrian behavior and precinct-level turnout.

Second, the data also has innumerable policy applications aimed at improving the safety and navigability of cities. By overlaying foot traffic data with real-time crime reports, we can improve our understanding of patterns of crime, making cities secure by implementing more efficient policing strategies and improving outdoor lighting and pedestrian safety. By layering data on transit infrastructure and use, we can better plan and build more accessible cities and improve emergency preparedness.

Finally, this technology allows for the real-time measurement of wait times at polling places, as well as counts of potential voters who "balk", or decide not to join the queue after seeing the length of the line. Such knowledge is crucial to improving election administration to ensure fair access to the democratic process, particularly with respect to racial minorities, who tend to face longer wait times (Pettigrew forthcoming; Alvarez et al. 2009). With a few exceptions, studies of wait times tend to rely on self-reports, which are notoriously subject to bias. Moreover, with the exception of Spencer and Markovits (2010), balking rates at polling stations and the relationship between balking and line length have not been systematically studied. Sensor-generated foot traffic data is revolutionary in that it provides real-time, objective measures of voting queue behavior that is not subject to human error. Thus it has the potential to improve the democratic process by removing a troubling yet persistent barrier to participation.

2.6 Conclusion

Cities tend to function as laboratories for new technologies that supply governments, residents, and businesses with better information about the pulse of the urban landscape. Although the data utilized in the above applications covers only a subset of one major American city, foot traffic data is becoming increasingly available to researchers. Placemeter's network of sensors is growing and the company plans to expand to other urban locales. Moreover, a number of private companies have invested in technology that generates foot traffic data by monitoring people's mobile devices; for instance, Google recently began rolling out a feature that uses such data.[17] Providing social scientists with access to these new data sources can facilitate research that augments our understanding of human behavior in ways we never thought possible, in turn leading to improved planning and policies.

[17] http://techcrunch.com/2015/07/28/google-search-now-shows-you-when-local-businesses-are-busiest/

3 Can violent protest change local policy support? Evidence from the aftermath of the 1992 Los Angeles riot (with Ryan D. Enos and Aaron R. Kaufman)

3.1 Introduction

Riots are political acts in which participants undertake violence to express grievances and attempt to spur policy change. Scholars and journalists often claim that riots cause short and long-term changes in political mobilization and attitudes and behaviors, both among riot participants and those exposed to the riot as observers or victims. Classic survey evidence has demonstrated that rioters ascribe political motivations to their actions (Sears and McConahay 1973), and scholars have argued that the series of riots in the 1960s caused shifts in the policy mood of the American electorate that had long-term consequences for national politics (Rieder 1985; Massey and Denton 1988; Edsall and Edsall 1992; Olzak, Shanahan, and McEneaney 1996; Western 2006; Manza and Uggen 2006; Wasow 2016). Similar claims have also been made about political violence in other countries (De Waal 2005; Hayes and McAllister 2001; Beber, Roessler, and Scacco 2014). In the wake of recent violent protest in African American communities in the United States and around the world, the subject of political violence has renewed importance.[1][2]

[1] We use the terms "riot," "violent protest," and "political violence" to describe large-scale violent protest in general, and the particular event in Los Angeles from April 29 to May 3, 1992. This is done for simplicity and without making assumptions about the motivations of the participants. We acknowledge that other writers prefer "uprising" and similar terminology.

[2] Replication data and materials are available on the Harvard Dataverse at Sands (2017a).

But can riots actually change policy? Measuring the effect of a riot on the policies said to motivate the rioters has been, as of yet, impossible. Despite the long literature examining riots and other protests as political acts, we have little understanding of whether these acts can achieve the goals ascribed to the rioters. We address this question by examining local shifts in referendum voting on goods targeted at urban-dwelling racial minorities after the 1992 Los Angeles riots. We find that the riot caused a shift in support for allocating these goods and that much of the shift is attributable to changes in mobilization among both African Americans and whites. We also find that this mobilization persists over time, evidence that the riot's political consequences were long-term.

In order to identify the effect of the riot on support for local referenda, we need not assume the underlying causes of the riot are unrelated to downstream changes in policy support. Rather, we assume the riot was exogenously timed relative to events planned long before the riot, namely an upcoming primary election. The triggering event for the 1992 Los Angeles riot was a video recording of police brutality and the subsequent acquittal of the police officers involved. Although the riot was a non-random event, the timing of the video recording and subsequent trial and conviction were exogenous to the timing of a primary election, thus granting causal leverage not found in most previous studies of political violence. We interrogate the validity of this causal claim by examining the spatial correlation between changes in policy support and the epicenter of the riot. We find that these changes were correlated with distance from the riot and not from other locations in Los Angeles, including other African American population centers.

3.2 Political Violence and Public Policy

By examining how a riot affected public goods allocation, our study addresses the effectiveness of rioting as a political activity. In doing so, we broaden our understanding of how political activity can bring about policy change, a question that typically examines more common political behaviors, such as voting, lobbying, or non-violent protest. While such activities are often argued to be ineffective in changing policy, especially for low-status groups (Gilens 2012) such as poor African Americans, we show that violent protest is an arguably efficacious political activity.

Of course, not everyone views rioting as politically motivated and not all riot participants are similarly motivated. However, social scientists often view rioting as a political act (e.g., Huntington 1968) and classic studies of riots show that participant aims include demanding redress for political grievances (Sears and McConahay 1973). Scholarly examinations of the 1992 Los Angeles riot also interpret the event as collective action against poor economic and social conditions (Tierney 1994). Indeed, post-riot polls indicated that 67.5% of African Americans in Los Angeles County viewed the riots as a protest against unfair conditions (Bobo et al. 1994). In addition to immediate grievances such as police brutality, riots are often characterized as demands for policy reform (Fogelson and Hill 1968). In the United States, these policies include spending on public goods benefiting the urban poor, including education, housing, and poverty assistance.

If participants riot in pursuit of policy, then understanding how riots affect policy support is important to understanding the effectiveness of this political tactic. Previous studies suggest that riots decrease support for racially liberal policies (Sears and McConahay 1973; Wasow 2016) or, conversely, that riots both increase negative attitudes toward the rioting group, and increase support for the policies advocated by riot participants (Beber, Roessler, and Scacco 2014). But because previous scholarship was limited to measuring post-treatment outcomes, considerable uncertainty remains about whether the riots actually caused these changes.

We concentrate on the effects of rioting on a local population. While the riot had national prominence, certain effects of the riot may have been localized—after all, a riot is a locally destructive event. Individuals close to the riot were more likely to be materially and psychologically impacted by the event and, as such, to be more motivated in the aftermath. Few studies have examined the localized effectiveness of a riot, instead treating the events as part of a larger phenomenon (e.g., Western 2006). While a broader focus can be valuable, it might mask differences between local and distal effects: for example, while the consensus in the literature that a riot makes citizens unsympathetic to the rioters may be accurate on a national scale, local opinion may become more sympathetic because of shared identity, special knowledge of local circumstances, or fear of further unrest. Additionally, examining the effects of a riot on national-level opinion, while potentially important, overlooks that much policy, especially in the domains of welfare and education, is controlled at the local or state level.

3.3 The 1992 Los Angeles Riot

On April 29, 1992, a court acquitted four white police officers for assault in the beating, captured on video, of an African-American man named Rodney King. The trial, relocated from Los Angeles to homogeneously-white Ventura County, concluded with an all-white jury acquitting the four officers of all charges.

Within hours of the verdict's announcement, a series of violent and destructive incidents occurred around the intersection of Florence Avenue and Normandie Avenue in south central Los Angeles, a predominantly African-American neighborhood. Reminiscent of the Watts Riots thirty years earlier, police officers abandoned the area, leaving residents to defend themselves against looters, arsonists, and widespread violence. For the next three days, the area suffered freeway shutdowns, suspension of municipal services, racial violence, and destruction of property. The violence proceeded unchecked until May 3, when 3,500 federal troops arrived to supplement 10,000 members of the National Guard. On May 27, one week before the 1992 primary election, the last troops withdrew from the area. The Los Angeles Riot resulted in 53 deaths, more than 2,300 injuries, and arrests in excess of 11,000; estimates of material losses exceed $1 billion (CNN 2013).

Much of the violence was covered live on television by news helicopters, including the beating and attempted murder of white truck-driver Reginald Denny as he attempted to drive through the intersection of Florence and Normadie. Contemporary accounts of the events describe great anxiety among the white residents of Los Angeles. In the aftermath came a flurry of reporting, often backed by survey data, speculating on the riot's effects on public opinion. Major themes in the media included an increase in fear, and a recognition that the conditions in which urban minorities lived needed to change (Toner 1992).

This initial view of aggregate opinion suggests that the riot represented a profound experience for Angelenos. Whether and how those experiences translated into policy support is our focus. Despite the fear and destruction caused by the riot, less than a month later voters in Los Angeles went to the polls in a statewide primary for electoral offices and ballot propositions. To understand the effects of the riot on voting in this election and in subsequent elections, we exploit the timing of this event to study changes in policy support.

3.4 The Local Effects of Violent Protest

In order to test for the localized effect of a riot on policy, we concentrate our study in the area of Los Angeles County closest to the riot. The impact of the riot is likely heterogeneous, with the greatest material and psychological impact on those closest to the riot. Unlike previous studies of riots that rely on aggregate units, we exploit geocoded individual- and precinct-level data to measure the effects of the riot on voters at any proximity to its epicenter, including those who lived at very close range. Using individual voters and precincts avoids problems of aggregation that are common to this type of analysis, such as the modifiable areal unit problem and problems of scale.

Los Angeles County is very large— spanning an area the roughly the size of Connecticut, it is the most populous county in the United States. As such, treating the riot as homogeneous in this area may hide important variation. As a principled way of defining "local" so as to avoid making arbitrary choices about the relevant distance from the riot, we focus on voters who are not separated from the riot by topographic features that may lower the salience of the riot. We concentrate our study in the area known as the Los Angeles basin, consisting of parts of the city of Los Angeles and other municipalities.[3] With this we exclude areas of the county further away from the riot, such as the desert communities to the east, and from areas that were separated from the riots by physical barriers, such as residents in the mountainous areas of the county or of the San Fernando Valley, who were separated from the riots by the Santa Monica Mountains. All data we describe below is limited to this area, which is entirely within thirty kilometers of Florence and Normandie, the epicenter of the riot.

We concentrate our study on two groups: non-Hispanic whites and African Americans.[4] These groups had substantially different baselines in terms of pre-riot ideology and political involvement, as well as divergent on-the-ground and psychological experiences with the riot. African Americans, whether riot participants or not, were much more likely than whites to share a social identity with the rioters and, perhaps, to sympathize with their grievances. Whites were *a priori* less likely to share a social identity with the rioters or to agree with their actions. In a nationally representative survey fielded during the riot, when asked whether the violence of the riot "was justified by the anger that blacks in Los Angeles felt over the verdict in the trial" only 17% of white respondents said it was, while 35% of African Americans did so.[5] Furthermore, the policy attitudes of the white population, compared to the African American population, were likely more heterogeneous, including more political conservatives: prior to the riot, whites, on average, would have been less likely to share the policy demands of the rioters than were African Americans.

We attribute differences in vote outcomes and other behavior before and after to the effect of the riot. This is necessarily a "bundled treatment" because a riot is associated with changes on a number of fronts, including psychological effects, media coverage, action

[3] We define a polygon that is south of the Santa Monica Mountains, extends east to the area of the city of Whittier and the small mountain range of Hacienda Heights and La Habra Heights, extends west to the Pacific Ocean, and south to Orange County. In our analysis we use precincts that intersect with this polygon.

[4] Hispanics, Asians, and other groups are excluded because of insufficient data, and because the theoretical expectation for these groups is less clear.

[5] Yankelovitch/Time Magazine/CNN Poll, April 30, 1992.

by politicians, changes in property values, reactions to the verdict, and many others. The purpose of this study is not to isolate the effect of each of these treatments, but rather to speak to the overall effects of a riot. As a large-scale collective political act, a riot is not unlike other such large-scale political acts, such as campaigns, where scholars have measured overall effects without identifying precise mechanisms. As such, we use the phrase "effect of the riot" as short-hand for the bundled effect of all the treatments associated with the violent protest.

To estimate the effect of the riot, we use differences in vote outcomes between June 1990 and June 1992. Other events associated with, but prior to, the riot, such as the beating of Rodney King or the subsequent trial, also occurred during this period and could affect policy support as well. We assume, however, that these events did not affect voter behavior nearly as much as the riot itself: there is no literature indicating that these types of events, even when racialized, can dramatically change voter behavior. With these assumptions in mind, in our empirical analysis we show that the effects we measure are likely attributable to the riot itself, not other associated events.

Are riots effective in bringing about policy change? In gauging the influence of a riot, researchers could examine survey data or the behavior of legislators and other policy makers. However, surveys may simply capture "cheap talk" or unstable attitudes (Zaller 1992) and the actions of legislators only indirectly reflect public opinion. As such, we directly examine local changes in mass-public voting on resource allocation to public schools.

We focus on support for spending on public schools as a public good which is associated with African Americans and racial minorities more generally and is often implicated in the social welfare demands of riot participants. Public schools have long been part of policy debates over how to address urban social and economic problems, including in official reports in the aftermath of urban rioting (California Governor's Commission on the Los Angeles Riots 1963). Research has demonstrated that demographic considerations play a role in opinions about public school spending (Hopkins 2009), with the non-white composition of schools affecting white support for spending. In 1992, on the eve of the riot, the Los Angeles

Unified School District was 72% non-Anglo-white. Near the riots, public schools were even more non-white, with enrollment at the 20 schools closest to Florence and Normandie 55% African American and less than 1% white.[6]

Our inferential strategy holds constant general attitudes about educational spending not associated with African Americans by comparing changes in support for public school spending to changes in support for university spending. Our assumption in doing so is that attitudes about university spending are not plausibly linked to African Americans and would not be as widely seen as a method of addressing problems made apparent by the riot. Thus, when white voters were asked to cast ballots with the demands of the African Americans rioters fresh in their minds, these demands would have no effect on votes on university spending. We believe this assumption is plausible: in our review of the literature, we found no scholarship claiming that attitudes about university spending were associated with attitudes about race or racial demographics.

The riot may have changed local vote outcomes through several channels, including mobilization, a general change in ideology or political outlook, or by altering the considerations citizens use when voting. The latter refers to the fact that most citizens do not have well-developed policy attitudes and so, when citizens form opinions, they are influenced by recent salient cues (Zaller 1992; Sands 2017e) which have been shown to influence vote choice (Berger, Meredith, and Wheeler 2008b). For citizens voting on questions of public policy so soon after the riot, the riot itself was likely one of these cues.

Effect on White Policy Voting

When the riot made salient to whites the policy demands of the rioters specifically and African Americans more generally, in what direction would this move their support for spending on public schools? We might expect a riot to increase support for spending by making white voters more aware of the needs of African Americans and the demands

[6] Those numbers were largely unmoved after the riot, with the percent African American dropping by 1.5% in the 1992/1993 academic year `http://www.cde.ca.gov/ds/sd/sd/filesenr.asp`

of the rioters. In particular, the signal generated by rioting—an extreme behavior—may have caused them to update their opinions about the severity of the needs of the African American community. The riot may have also activated pre-existing sympathy for these policy positions. Or, the destructiveness of the riot may have caused white voters to favor investment in programs that would prevent further violence (Beber, Roessler, and Scacco 2014).

Among elites, there is a history of such reactions to riots in the United States. The Kerner Commission, formed by Lyndon Johnson, attributed the urban unrest of the late 1960s to a host of inadequate welfare institutions. The commission recommended reforms of social welfare programs to prevent further riots (United States National Advisory Commission on Civil Disorders 1968). A report after the 1963 Watts Riots by a commission chaired by former CIA Director John McCone made similar recommendations on reforms, including for improved schooling for African American children (California Governor's Commission on the Los Angeles Riots 1963). After the 1992 riot, there is evidence that voters nationwide drew similar lessons. A series of post-riot surveys found that 65% of respondents agreed that "the violence in Los Angeles…has made it more urgent to address poverty".[7] and that 51% wanted to see an increased emphasis on providing social services, as compared to 37% who supported an increased emphasis on promoting law and order.[8] Moreover, 38% of respondents desired more spending for minorities in urban areas, while only 13% preferred decreased spending.[9] Perhaps the relatively liberal voters of the Los Angeles basin were moved by the logic of needing improved social institutions and, thus, more spending to prevent future riots.

However, when asked to make consequential decisions about spending rather than merely state preferences on a survey, would white voters of Los Angeles County, where the riot had the greatest impact, have the same liberal reaction? The literature offers reason to predict

---

[7] Time, Cable News Network Poll. , May 13–14, 1992.

[8] NBC News, Wall Street Journal Poll. May 15–19, 1992.

[9] Gallup Organization Poll, May 7 - May 10, 1992.

that the riot would dissuade voters from spending on public goods associated with a racial outgroup.

For example, it is well-established that individuals display ingroup bias in attitudes, behaviors, and resource allocation, all else equal (Fiske 2000). This behavior, repeatedly observed in the laboratory, is often cited as one of the primary reasons for the well-established finding that cross-nationally, diversity is negatively correlated with public goods provision (Alesina and Zhuravskaya 2008; Habyarimana et al. 2009; Enos and Gidron 2016). Scholars have drawn direct connections between local ethnic diversity and the willingness of voters to allocate funds for public goods, including schools (Alesina, Baqir, and Easterly 1999; Rugh and Trounstine 2011).

Furthermore, research has shown that whites in the United States draw on negative stereotypes about African Americans when forming policy attitudes about benefits perceived to disproportionately target African Americans (Gilens 1999). In laboratory experiments, making race or other group-based identities salient leads to group-based discriminatory attitudes and behaviors (Sidanius and Pratto 2001; Enos and Celaya 2015; Reicher et al. 2016). With negative stereotypes made salient by the televised images of violence and looting, it is possible that whites drew on these stereotypes when asked to make decisions about allocating to public schools. With this substantial literature in mind, we might also predict that the riot, which brought attitudes about the outgroup to bear on voting decisions, would precipitate a decrease in support for those public goods.

Effects on African American Policy Voting

African Americans' perspective on the riot was likely divergent from that of white voters: a group with whom they share a racial identity participated in widespread violent protest. Given the strong connection between group identity and policy attitudes, grounded in shared history (Dawson 1995), the riot may have increased African Americans' awareness of ingroup policy demands, thereby causing a shift in attitudes.

Salient political events, such as electoral campaigns, can help citizens learn about the policy preferences of groups with which they share an identity. This learning process can bring attitudes into line with the dominant attitude of the group (Lenz 2013). This might occur because voters use group membership as a heuristic for the "right attitude", which they adopt upon learning about the group's preferences (Cohen 2003). The salient politicized event of a riot could serve the same function as a campaign: the policy demands of the rioters are extensively covered by the media and likely transmitted through social networks, thus the attitudes of people who identify with the rioters will change to reflect these demands. Of course, African Americans in Los Angeles may have already had largely liberal policy preferences and attitudes that were crystallized before the riot; if so, the riot itself may have little or no effect on their attitudes. However, we do not see how the riot could cause a *decrease* in support for spending among African Americans.

3.5 Data and Estimation Procedure

We test for the effect of the riot on support for policies favored by the rioters by examining voting on ballot initiatives related to funding for public schools. To isolate the effect of the riot, we examine votes on two related sets of public goods: public schools, which were closely linked to African Americans and may have been seen as a way to address to policy concerns of the rioters, and universities, which were not linked to African Americans and were, at most, a very indirect way to address the policy concerns of the rioters.

We focus on four initiatives that appeared on the ballot in the June election in 1990 or 1992: Prop 121 (1990), Prop 123 (1990), Prop 152 (1992), and Prop 153 (1992). These propositions are summarized in Table 3.1. We rely on the symmetry of ballot initiatives from 1990 to 1992 for our test of the effect of the riot: both elections have both a public schools and a university education initiative and the monetary amounts associated with both initiatives in 1992 are approximately twice what they were in 1990.

We measure a "difference-in-differences" from these four ballot initiatives:

$$\texttt{EdDiff}_i = (\texttt{PubSchool}_{i1992} - \texttt{PubSchool}_{i1990}) - (\texttt{HigherED}_{i1992} - \texttt{HigherED}_{i1990})$$

Initiative	Year	Title	Dollar Amount	Statewide %	LA County %
121	1990	Bonds for Higher Education Facilities	\$450 M	55.0%	59.7%
123	1990	Public School Construction Bonds	\$800 M	57.5%	60.3%
152	1992	Bonds for Public Schools	\$1.9 B	52.9%	57.7%
153	1992	Construction Bonds for Higher Education	\$900 M	50.8%	56.8%

Table 3.1: Summary of ballot initiatives

$\texttt{EdDiff}_i$ is the change in support for public schools in precinct i between 1990 and 1992, net the change in support for universities. Our primary quantities of interest are the population-weighted mean of $\texttt{EdDiff}_i$ for all voters, and for white and African American voters separately. By "differencing out" the change in support for universities from the change in support for public schools we substantially reduce the threat of omitted variable bias. Our estimate of the effect of the riot on public school spending will be biased if there was a shock to support for university funding between 1990 and 1992 that was not also a shock to public school funding; however, we know of nothing that would cause such a shock during this time period. Because our estimation strategy nets out any secular change in, for example, voter ideology, we believe we have isolated the change in support for public schools that is due to the riot. To further isolate the effect of the riot, we also check for correlations between $\texttt{EdDiff}_i$ and distance from the riot. As noted below, significant correlations between these variables give us confidence that we are isolating the riot's effect; for an omitted variable to bias our estimate it would have to be correlated with both distance and $\texttt{EdDiff}_i$.

To acquire precinct-level data, we digitized vote returns from the Los Angeles County Registrar-Recorder County Clerk's Elections Division. This data set allows us to examine the aggregate effects of the riot. However, while we know the total vote outcomes and the racial demographics for each precinct, to understand the distinct voting patterns of different racial groups requires individual-level vote data. This presents an "Ecological Inference Problem" (Robinson 1950). We use the Ecological Inference methods developed by King (1997) to isolate behavior by racial group. While, as with all estimation techniques, this method relies on assumptions for validity, the technique has been validated on well-known problems

(O'Loughlin 2000) and used for inference in other important questions (Enos 2016).

The inputs to the EI model are: the proportion of whites, African Americans, Hispanics, Asians, and others in each precinct; and the proportion of Yes and No votes for each ballot initiative. The outputs are estimates of what proportion of each group voted Yes on a given ballot initiative, for each precinct in the data set. With these outputs, we separately measure $\texttt{EdDiff}_i$ for whites and African Americans using the equation above. Our measures of precinct-level demographics, including the racial characteristics of each precinct, come from the California Statewide Database, which merges local voter files and decennial U.S. Census data with precinct geographies to create demographic counts by precinct.

To explore mobilization, we use data from the 1992 Los Angeles County voter file, which includes 3,743,468 registered individuals, and contains address, sex, age, and party registration information for each person. Using ArcGIS, we geocode voter file addresses to match individuals to spatial coordinates, then place those coordinates inside precincts, which allows us to measure the distance of voters and precincts from the riot and to validate our claim that changes in policy support were caused by the riot. We measure the geospatial distance between the center of the riot at Florence and Normandie and the interior centroid of the voting precincts used in our analysis. Because precinct-level GIS data from this time-period do not exist, to create this data, we geocoded every address from the 1992 Los Angeles County voter file[10], and calculated population-weighted precinct centroids from the addresses of every voter in the precinct. To validate these geocodes, a team of research assistants used Google Earth to digitize precinct maps.

We supplement the voter file with imputed information about the voters' race. We merge the voter file with two additional data sets: Census Block-level race and ethnicity data, and the Census surname list, which indicates the probability of belonging to each race given a surname. Using the surname probability as a prior, we apply Bayes' Rule to compute the posterior probability that an individual is White, Hispanic, African-American, Asian,

[10] Of the 3.7 million addresses in the voter file we successfully geocoded 99.62% in ArcGIS.

Native American, or other, conditioning on their address and surname (Enos 2016).[11] From this vector of posterior probabilities, we identify the race with highest probability and use it as the imputed race of the voter. Such imputations have been used successfully in prior research and are likely quite accurate in a segregated area like the Los Angeles basin.

Mobilization Data

From the 1992 voter file, we measure the change in partisan registration between the period just before to the period just after the riot. We do this by comparing a one week window on either side of the riot, which lasted from April 29 to May 3. This means we compare a pre-riot period of April 22 to April 28 and a post-riot week of May 4 to May 8. In the pre-riot week, we omit the Saturday and Sunday of April 25 and 26. Note that these periods contain the same days of the week, which is advantageous because registration tends to vary by day of the week.

However, if voters are more likely to register in weeks closer to the registration deadline, our pre-post comparison may be confounded. The coincidence of the timing of the riot vis-à-vis the voter registration deadline of May 4 allows us to overcome this problem. Because the riot occurred in the days leading up to the registration deadline and prevented people from registering, election officials announced after the riot that they would extend the registration deadline by three days. This provides a unique opportunity because it means there are two roughly equal periods of time prior to the registration deadline. Prior to the riot, people had no way of knowing that the deadline would be extended and thus would register as normal. After the riot, a new deadline appeared and everyone who was motivated to register because of the riot was able to do so. In this way, we can capture the marginal effect of the riot on the propensity to register as a Democrat or Republican, while controlling for seasonal and

[11] We were unable to create a geographic prior for a small portion of the names on the voter file. This was either because a slight imprecision in geocoding placed the addresses in a zero population Census Block (for example, an all-commercial block), or because we were unable to geocode the address to begin with. In these cases, we rely exclusively on surname data to impute race. Similarly, individuals whose surnames do not appear in the census surname list have their race imputed using their census geography only. These two methods were used for 3.6% of the voter file. For no individual did we have neither surname- nor census-based racial information.

other unobserved trends that also affect registration.

Attitude Data

We use individual-level survey data from the Los Angeles County Social Survey (LACSS) to measure attitude change. The LACSS, conducted annually from 1992 to 1998 (University of California, Los Angeles. Institute for Social Science Research 2011), captured a county-representative sample generated by Random Digit Dialing. The 1992 LACSS, themed "Ethnic Antagonism in Los Angeles," provides a picture of racial attitudes and policy preferences expressed by a random sample of Angelenos before and after the riots. By coincidence, the verdict and the riot occurred in the middle of survey implementation, enabling us to leverage variation in survey responses recorded immediately before and after this exogenous shock, as done by Bobo et al. (1994).

Long-term Data

Finally, to examine long-term partisanship and participation, we merge the 1992 Los Angeles County voter file with a statewide California voter file from 2005. This was the earliest voter file we could obtain that included voter turnout data. By examining names, gender, and date of birth of the 30,166 voters who registered in the 10 weekdays before and after the riot, we successfully locate 15,244 of them in the 2005 file. Our calculations suggest that after accounting for voters who moved out of state, died or changed their name, we successfully locate a high proportion of our sample who were still registered in 2005. We detail this matching process in the Appendix.

3.5.1 Results: Changes in Policy Support

We first consider the aggregate change in Angelenos' support for public school spending as a result of the riot. If, on average, the riot served as a negative shock to support for public school funding, then the mean of $\texttt{EdDiff}_i$ should be negative. If the riot served as positive shock for support public school funding then the mean of $\texttt{EdDiff}_i$ should be positive. These results do not rely on ecological estimates and, therefore, there is no model-based uncertainty around the result.

The distribution of $\texttt{EdDiff}_i$ is displayed in the top panel of Figure 3.1. The vertical dotted line is at the mean of the distribution. Pooling together all precincts in the Los Angeles basin, we see a population-weighted mean $\texttt{EdDiff}$ of 0.052 (95% confidence interval: $[0.048, 0.055]$)[12], indicating that average support for public schools, net of the change in support for higher education, increased. The difference in support for university funding between 1990 and 1992 was -0.021 (CI : $[-0.025, -0.017]$), clearly showing that overall support for education spending did not increase, but rather this change in support was specific to local public schools.

Voters proximate to the riots experienced a large positive shift in their support for public schools, even when accounting for their overall shift in support for education. Because public schools are a public good closely associated with the issues related to the riot and to the social identity of the rioters (as African Americans), this provides initial evidence that the riot was effective in generating support for the rioters' policy demands.

Comparing these results to the same difference-in-differences in other parts of California bolsters the claim that this effect is due to the riot. The statewide difference-in-differences was close to zero, and the county-level difference-in-differences in other large, urban counties not experiencing the riot, such as San Diego County (-0.007) and San Francisco County (-0.009), were also close to zero.

Results by Race

The population weighted means for both whites (0.033, CI: $[0.029, 0.036]$) and, especially, African Americans (0.080, [CI: $0.077, 0.083$]) demonstrate an increased willingness to pay for public schools relative to universities (see bottom panels of Figure 3.1). Also evident in Figure 3.1, the distribution of $\texttt{EdDiff}_i$ is wider for whites than it is for African-Americans, reflecting more variation in responses to the riot, as might be expected given more pre-riot

[12] All confidence intervals in this manuscript are bootstrapped.

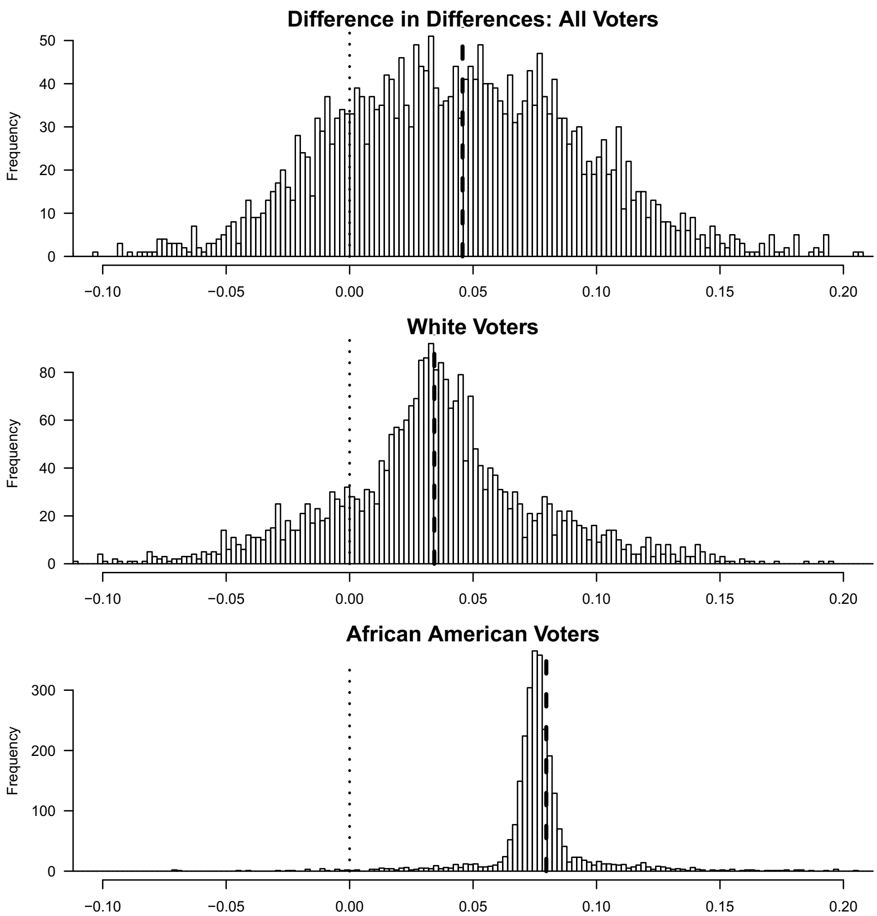

Figure 3.1: Histograms represent distribution of $\texttt{EdDiff}_i$ for all voters (top), whites (middle), and African Americans (bottom) at the precinct level. Positive values represent an increase in support for public schools, net of changes in support for universities. The vertical line is the population-weighted mean of the difference-in-differences.

heterogeneity among white voters in terms of distributive preferences.[13] The larger and more uniform shift by African Americans compared to whites is consistent with the claim that these changes are responses to the riot and not something else, and suggests that a significant portion of the effectiveness of violent protest comes from rallying support from people sharing an identity with the rioters, rather than from gaining support from outside groups.

[13] A weighted T-test of the differences in means between `EdDiff` for whites and African Americans, yields $t = -38.75$, corresponding to a p-value < 0.01.

This initial evidence establishes that a riot can be locally effective in driving policy support, whereby members of the public, even those perceived to be the target of the riot (whites) and especially those sharing an identity with the rioters (African Americans), vote in a way that meets the policy demands of rioters.

3.5.2 Changes in Policy Support by Distance from the Riot

We now turn to tests to better establish that the changes we observe were caused by the riot itself and are not a spurious association. Because we lack sufficient data to show parallel trends— the standard validity check in difference-in-differences analyses— we turn instead to correlations with distance to buttress our causal claims.

If changes in policy support were caused by the riot, then we would expect these changes to be correlated with distance from the riot because the salience of events can vary with a subject's proximity to the event (Latané 1981). Note that the direction of this effect may cut both ways; changes in support for public schools may increase or decrease with distance. Since the overall effect of the riot was an increase in support for school spending, we might expect this increase to be largest nearer to the center of the riots. However, there could be psychological forces that diminish as proximity to riot increases: a literature from psychology has demonstrated that traumatic events, such as violence, can induce increased discrimination toward the outgroup because of basic psychological motivations to look to the ingroup for protection during threat (Navarrete and Fessler 2005) or because fear induces a need for belief-reinforcing self-esteem (Greenberg, Pyszczynski, and Solomon 1986). In fact, medical accounts of witnesses to the riot describe severe emotional trauma, and symptoms resembling post-traumatic stress disorder (Kim-Goh et al. 1995). Drawing on this literature, we might expect that white individuals for whom proximity to the riot induced greater fear will be more likely to withhold public goods from the outgroup than those living further away.

To test the relationship between distance and `EdDiff`, we estimate population-weighted precinct-level regressions treating $\texttt{EdDiff}_i$ as our dependent variable and distance as our

	White	Black	White	Black
Distance	-0.006	-0.011^{*}	0.063^{*}	-0.021
	(0.004)	(0.004)	(0.022)	(0.017)
Distance^2			-0.028^{*}	0.006
			(0.009)	(0.009)
Constant	0.041^{*}	0.085^{*}	0.001	0.088^{*}
	(0.006)	(0.003)	(0.013)	(0.006)
Observations	1,534	1,534	1,534	1,534
R^2	0.002	0.006	0.010	0.006

Table 3.2: Relationship between `EdDiff` and distance from Florence and Normandie. OLS regression of $\texttt{EdDiff}_{\texttt{i}}$ and distance from Florence and Normandie, measured in units of 10 kilometers. Standard errors are in parentheses. $*$ represents $p < .05$.

explanatory variable. These results are in Table 3.2. We include both linear and quadratic terms on distance to account for potentially nonlinear effects and find that as distance from Florence and Normandie increases, whites are more likely to support public schools relative to universities. In other words, whites closest to the riots experienced less change in their support for public school spending compared to whites further away.

In substantive terms, the quadratic model in the third column of Table 3.2 shows that a white voter residing in a precinct located 1 km from the epicenter of the riots has a predicted `EdDiff` of 0.006, a relatively small change in support for public schools relative to universities. Meanwhile, a white voter located 10 km from the riot has a positive predicted `EdDiff` of 0.035, suggesting a large increase in likelihood of supporting public schools relative to universities.[14] That the riot increased support for public schools, a public good disproportionately serving African Americans, speaks to the effectiveness of the riot in achieving policy change. However, that this support declined for those closest to the riot speaks to a potential for reduced effectiveness of political violence to garner support if the fear and

[14] One potential explanation for this is that the pre-riot levels of support near the riot center were so high that this would present a "ceiling effect" whereby there was no room to increase. However, this appears not to be the case. Precincts within 1km of Florence and Normandie had 1990 levels of support at 57.7% for the universities initiative and 55.6% for the public schools initiative. These were close to the averages for precincts further away.

damage it engenders counteracts increases in sympathy.

By comparison, the relationship between distance and `EdDiff` for African American voters is substantially weaker and appears to be in the opposite direction of the relationship for whites. The relationship between distance and African American voter behavior is better captured by the linear model in column 2. For this group, the estimated coefficient on distance is negative and statistically significant, but relatively small in magnitude. For a African American voter residing 1 km from the riot epicenter, the predicted value of `EdDiff` is 0.084; her counterpart living 10 km away has a predicted value of 0.074. This is further evidence that African Americans experienced a more uniform change in support that cut across income and demographic lines, as might be expected if African Americans widely identified with the rioters.

That policy support change was correlated with distance from the riot reinforces the claim that the observed shifts were in response to the riot and not to some spurious variable. In order for an observed or unobserved variable to threaten the validity of our analysis it must be correlated with *both* `EdDiff` and with distance from the riots. In other words, any potential confounder would need to correlate with the difference in the change in voting over time and also with our measure of distance. Furthermore, this is evidence that the effect we measure in `EdDiff` is the effect of the riot itself and not of other related events, such as the media coverage of the trial or the beating, because there is no reason to believe changes in support for public goods caused by those events would be correlated with distance from the location of the riot.

Changes in Policy Support by Distance from other Places in Los Angeles

As a placebo test, we check for a systematic relationship between `EdDiff` and distance from locations other than Florence and Normandie. If there are no such relationships, this gives us greater confidence that the change is a reaction to the riot and not an unobserved variable. For each voting precinct in our data, we calculate the distance to that precinct from every other precinct, and, for white voters, regress `EdDiff` on distance using these alternative distance measures. We then compare the magnitudes of the resulting coefficient estimates to those described in Table 3.2.

We plot the results of this placebo procedure in Figure 3.2. Areas are shaded by the size of the African American population and the white star indicates the location of Florence and Normandie. Each point represents a placebo test, using distance from that location as the independent variable in the regression on `EdDiff`. The largest points correspond to places where the coefficient on distance is similar to that described in our original analysis above; smaller points indicate smaller coefficients. Note that the only area of statistically significant coefficients is located immediately south and west of Florence and Normandie; similarly, placebo tests conducted using distance from the large African American populations to the northwest produce coefficient estimates close to zero. Overall there is no evidence of a similar relationship between changes in policy support among whites and distance from any other point, other than those also near the riot. In the Appendix we show results from an additional placebo test, examining the relationship between distance and the trend in ballot measure support in the years prior to the riots.

3.6 Exploring Possible Mechanisms for Increases in Policy Support

We have shown that Los Angeles voters increased their support for spending on public schools after the riot, holding constant support for spending on education more generally. Public education is a policy area associated with urban poverty and implicated in the demands of the rioters. This change, we claim, was a reaction to the riot, implying that violent protest can be a locally effective policy tool. We demonstrated that the change in support was much greater for African Americans than for whites, that changes in support were correlated with distance from the riot, and that changes in support were not correlated with distance from other locations.

We now turn to testing the mechanisms for this phenomenon by examining attitude

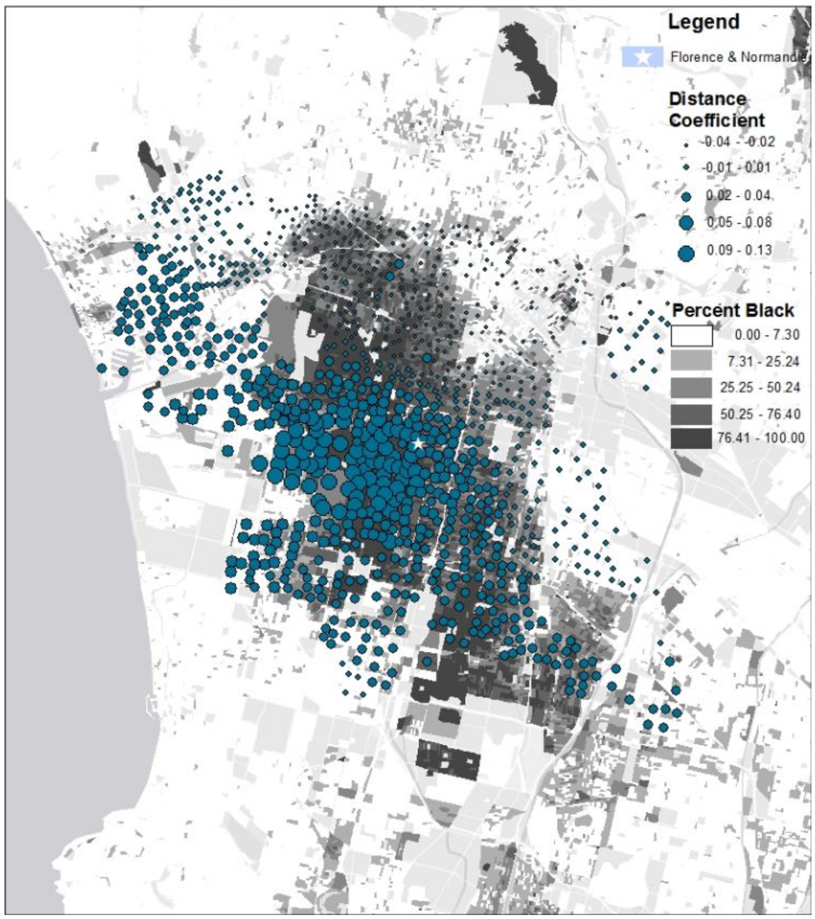

Figure 3.2: Relationship between distance and other points in the LA Basin. The coefficient on distance from other points and EdDiff. Larger points indicate higher (more positive) coefficients. Precincts are shaded by percent African American.

change, changes in the electorate, and changes in mobilization. While we cannot establish definitive connections between these variables and shifts in policy support, we believe the data more convincingly point to mobilization as the primary driver of the change.

3.6.1 Attitude Change

The LACSS data, the collection of which spanned the weeks surrounding the riot, allows us to test whether changes in aggregate policy support were accompanied by changes in attitudes among survey participants. The validity of our claims in this section rests on the assumption that those selected to be surveyed prior to the verdict are not systematically different from those sampled afterwards. To test this assumption, we examine balance on available covariates separately for whites and African Americans before and after the riots. The white sample is small ($n = 185$) and fails this balance test. As such, we do not examine changes in white policy attitudes. African American subjects, on the other hand, were well balanced across pre- and post-riot samples and comprised a sufficiently large sample ($n = 426$). Balance tables for both groups are displayed in the Appendix.

Regressions based on these survey results suggest that, despite changes in aggregate voting behavior, African American support for spending on "improving our nation's education system", as well as ideology on a liberal to conservative scale, remained constant post riot.

3.6.2 Changes in the Electorate

Another possible explanation for changes in policy support is that citizens exited the electorate due to the riot or to other, non-riot related, factors occurring at the time. Citizens may have moved, or simply "hunkered down" after the violence. Perhaps, for example, citizens opposed to spending on public education disproportionately moved out of the basin after the riot. Although difficult to test directly, we believe it is implausible that citizens relocated *en masse* within the brief intervening time period between the riot and the June election.

On the other hand, that citizens hunkered down and chose not to participate because of the riot seems more plausible. If this occurred, it would be evident in changes in voter turnout rather than registration rates. This is not immediately apparent in the data: turnout among registered voters in Los Angeles County increased by 10.6 percentage points between the 1990 and 1992 June elections. However, changes in turnout must be considered in light of secular changes that typically occur between a midterm election year, like 1990, and a presidential election year, like 1992. Because turnout will be higher in a presidential year, decreases in turnout caused by the riot may have been masked by this secular trend. However, this does not appear to be the case in Los Angeles County because the increase in turnout between 1990 and 1992 was unusually high: it was higher than the average change of 6.0 percentage points in the next five largest California counties.[15] This increase was also higher than the average change in turnout in the county between the two previous midterm and presidential primaries (2.1 percentage points) and the following two midterm and presidential primaries (7.6 percentage points). Contrary to a hunkering down mechanism, it appears that voter turnout did not decrease in Los Angeles because of the riot, and in fact may have increased dramatically.

3.6.3 Mobilization

Salient political events, such as elections, can cause voters to become active in politics through increased voter registration (Meredith 2009). Could citizens have become politically active because the riot—a salient political event—raised their interest in politics, or because politicians used the riot to mobilize voters?

The riot may have mobilized African American voters who saw members of their ingroup intensely involved in politics, serving as a cue for the importance of political involvement or allowing for network effects to activate political participation. Recent violent protests in Ferguson, Missouri and Baltimore, Maryland were followed by other forms of political participation in African American communities, such as continued protest via the Black

[15] In 1990, by number of registered voters, these were San Diego, Orange, Santa Clara, Alameda, and Sacramento Counties.

Lives Matter movement. However, riots are also described in the scholarly literature as a form of political action undertaken because other forms are ineffective or unavailable (Sears and McConahay 1973). Thus it is possible that African Americans, having turned to rioting because of frustration with more mainstream forms of participation, will not be mobilized by the riot into mainstream politics.

White voters may also be mobilized by the salient political event of the riot, but the ideological direction of the mobilization is less clear than for African Americans. Both conservative and liberal politicians may exploit the event to register sympathetic voters. While previous scholarship finds that African American political activism, especially when violent, leads to a conservative backlash in voting among whites, these studies do not explore voter registration and so it is unclear whether this backlash comes from newly mobilized voters. Alternately, low-propensity voters are also more likely to be Democrats (Enos, Fowler, and Vavreck 2014), so it may be previously unregistered Democrats who are most motivated by riot, causing an liberal uptick in registration.

Because we measure mobilization through changes in voter registration, we cannot separate political conversion, where voters who would have registered anyway register with a different party, from pure mobilization, where voters who would not otherwise have registered do so because of the riot. However, in the survey data we see little ideological change among African Americans and, as we show below, levels of new registration were unusually high after the riot. Thus we have some reason to believe our findings reflect "pure" mobilization rather than conversion.

The riot appears to have a significant mobilizing effect. Examining the voter file, we see that in the week after the riot, despite the passage of the original registration deadline, 24,587 Los Angeles basin residents registered to vote. In the 20 weeks leading up to the riot, the mean number of registrants per week was 3,614 and the median was 3,292. In the week immediately prior to the riot, with the registration deadline looming, 5,579 people registered. The riot appears to be a tremendous mobilizing force.

A question remains as to whether this spike in registration was simply due to pent-up

	Before Riot	After Riot
African American	0.14	0.28
White	0.66	0.54
Age	35.23	34.26
Male	0.52	0.53
Distance from Florence/Normandie	14,649	12,387
N	5,579	24,587

Table 3.3: Demographics of registrants immediately before and after the riot. Before Riot includes new registrants on the five weekdays before the riot (4/22–24 and 4/27–28) and After Riot includes the five weekdays after the riots (5/4–8). All variables are proportions, except for age (years) and distance (meters).

demand, i.e., a result of those who would have registered anyway but were prevented by the riot from doing so, or whether it reflects those who would not have otherwise registered in the absence of the riot. Isolating these two groups is impossible, but comparisons to other elections suggest that this spike reflects registration by those who would otherwise not have registered. The 2004 presidential primary in California provides a reasonable baseline because, as in 1992, the primary election involved choosing a Democratic candidate to run against an incumbent Republican president in the general election. Prior to the 2004 presidential primary, 31% of Angelenos who registered in the four weeks prior to the deadline did so in the last five weekdays leading up to the deadline. Meanwhile, those who registered in the five weekdays before the new, post-riot registration deadline in 1992 represent 45% of those who registered in the preceding four weeks.

Not only did the riot appear to be a mobilizing force, but it appears to have registered a different demographic than those registering in the absence of a riot. Using the covariates available in the voter file, we see systematic differences in the composition of those who registered just before and just after the riot. We summarize this data in Table 3.3. Notably, compared to those registering before the riot, those registered after the riot are of similar age, but are more likely to be African American, less likely to be white, more likely to be female, and more likely to live closer to Florence and Normandie (see Appendix). This too is consistent with a mobilizing effect of the riot.

We also examine whether the riot changed the probability that new voters register as

	Before riot	After riot	Difference	p value	N
White Pr(Reg Republican)	0.36	0.25	-0.11	0.00	3,392
White Pr(Reg Democratic)	0.64	0.75	0.11	0.00	9,083
Black Pr(Reg Republican)	0.04	0.03	-0.02	0.05	184
Black Pr(Reg Democratic)	0.96	0.97	0.02	0.05	6,104

Table 3.4: Post-riot partisan shift among new voter registrations by whites and African Americans. P values generated by T-test for difference of means.

Democrats or Republicans. We display these results in Table 3.4, where Pr(Reg Republican) is the proportion of registrants who affiliated with the Republican party out of the total new registrants, and Pr(Reg Democratic) is the proportion who registered as Democrats. Consistent with the overall liberalizing effect of the riot, both white and African American voters decreased Republican registration and increased Democratic registration; however, given the already very high Democratic identification among African Americans voters, the changes among this group are small.

Our ballot data indicates that the difference in votes cast for public schools and universities went from -7,798 votes in 1990 to 64,448 votes in 1992, a swing of 72,246 votes. Although it is impossible to determine how much of this swing was due to the mobilizing effect of the riot, we note that the 24,587 people registered in the wake of the riot could represent a large portion of that shift. Furthermore, most of the new registrants were African American Democrats, the voters who appear to have shifted the most in support of public schools.

As noted above, the riot may have increased mobilization through a variety of channels and, as such, is a bundled treatment. The increase could come from a piqued interest in politics that results directly from the riot or from the indirect effect of politicians using the riot to mobilize supporters. However, these shifts are particularly notable because such large-scale changes in party registration over this short of a period do not usually occur in modern American politics (e.g., Green, Palmquist, and Schickler (2002); Erikson, MacKuen, and Stimson (2002)), suggesting that the riot was a rare and dramatic political event.

3.7 Long-term Shifts in Registration and Participation

This mobilization may have downstream effects as those mobilized remain active in politics and influence policy at a later date. We check for the long-term persistence of both party identification and voter participation among those mobilized by the violent protest to measure the long-term effectiveness of violent protest as a policy tool. Voters are socialized by salient events which shape their long-term partisanship, so we might believe that voter registration resulting from the riot will show greater partisan stability than usual. However, the riot being an extraordinary event, possibly used by politicians to register voters who would otherwise not be registered, might mean that voter participation activated by the riot will wane over time.

Merging the 1992 voter file with the 2005 file allows us to observe the behavior of those treated by the riot, over a decade after it occurred. In Table 3.5, we show changes in party registration probabilities between 1992 and 2005 of both white and African American voters, among those who registered immediately before or after the riots. Both white and African American voters who registered after the riots were more likely to register as Democrats, and less likely to register as Republicans, in 1992. Looking at those same individuals in 2005, party registration is highly consistent, with those registering after the riot, regardless of party, as likely, if not more, than those registering before to have the same partisan affiliation in 2005 as in 1992.

Among white voters, 82% of Democrats and 77% of Republicans registering after the riot kept their party affiliation, comparable to those registering before the riot. For African Americans, 87% of Democrats registered after the riot remained Democrats. This is identical to the stability for Democrats before the riot and higher than the stability for whites. African American Republicans show considerably less stability, but their number is small.

We can benchmark this registration by examining the stability of partisanship reported in Green, Palmquist, and Schickler (2002), who argued that partisan identity is a social identity after comparing its long-term stability to other important identity-categories, such

White Before-Riot

		2005		
		Dem	Rep	Total
1992	Dem	0.82 (683)	0.08 (70)	0.90 (753)
	Rep	0.12 (52)	0.77 (327)	0.90 (379)

White After-Riot

		2005		
		Dem	Rep	Total
1992	Dem	0.81 (2993)	0.08 (292)	0.89 (3285)
	Rep	0.13 (145)	0.77 (831)	0.90 (976)

African American Before-Riot

		2005		
		Dem	Rep	Total
1992	Dem	0.87 (291)	0.03 (10)	0.90 (301)
	Rep	0.35 (6)	0.59 (10)	0.94 (16)

African American After-Riot

		2005		
		Dem	Rep	Total
1992	Dem	0.87 (2598)	0.04 (118)	0.91 (2716)
	Rep	0.36 (29)	0.52 (42)	0.88 (71)

Table 3.5: Long-term partisan stability for those registering pre and post riot. Proportion of party registrants in 2005 by registration in 1992, separately for whites and African Americans registering before and after the riot. Columns totals do not add up to 100% because of movement to minor parties or non-declared.

as race and religion. Using survey panel data, they found that between 1992 and 1996, 78% of partisans kept their same affiliation. In our data, those registering in Los Angeles in 1992 in the wake of the riot were at least, if not more, likely to keep their same partisan affiliation.

Having established that the extended registration period produced lasting effects on registrants' party affiliations, we turn to the question of long-term participation. In tables in the Appendix, we examine the difference in turnout in the 2004 primary and general elections for those registering before and after the riot by regressing turnout in those elections on the variable indicating whether an individual registered immediately before or after the riot. Notably, both with demographic controls and without, for both whites and African Americans over 10 years after the riot, the coefficient on registering before or after the riot does not attain statistical significance and, in several cases, is a precisely-estimated zero, indicating that those who were motivated to register by the riot voted just as regularly as those registering under normal circumstances, even when holding demographic differences constant. This is true even in primary elections where participation is often limited to the most politically active citizens. Combined with the stability of partisanship, the long-term mobilization of post-riot registrants as regular participators is a sign of the potential long-term effects of the riot: citizens appear to have joined the party sympathetic to the demands of the rioters and not only remained in that party, but remained active citizens. If the shift in policy support observed in the referendum voting remains in the long-term vote choices of these citizens, this speaks to the long-term ability of rioting to affect policy.

3.8 Conclusion and Discussion

We have provided evidence, which is plausibly causally identified, for the effect of a violent protest on local policy support. Our results indicate that a riot can help to accomplish policy or symbolic goals by mobilizing supporters or building sympathy among others. We showed that white and African American voters were mobilized to register as Democrats and shifted their policy support toward public schools, net of a general shift in support for education. This mobilization appears to have persisted: those mobilized by the riot remained regular participators over a decade later and remained more Democratic than the general population, even after accounting for demographics.

3.8.1 Why Do Our Results Differ from Previous Scholarship?

Our finding is strikingly inconsistent with previous literature. In both the American and comparative context scholars have found that political violence is associated with a "backlash" effect; voters behave unsympathetically towards the perceived rioting group. In the United States, this means increased support for social and politically conservative candidates and policies. A common argument in the literature is that the string of riots in the 1960s caused large proportions of white voters to abandon support for the liberal welfare state, which, especially since the Great Society of Lyndon Johnson, had been rhetorically framed around curing inner-city poverty associated with African Americans. Those riots are said to have caused changes in partisanship and voting behavior that ushered in the rise of Nixon and Reagan-era conservatism that still affects American politics today.

Our findings may differ from these prior studies because the series of riots in the 1960s had non-linear effects; perhaps while a single riot invokes sympathy, a series of riots provokes backlash. The difference between local and non-local effects may also be consequential: those observing the riot from afar may lack sympathy because they do not share an identity with the rioters. Much of the previous literature examining the political effects of urban riots in the United States neglects whether those riots helped galvanize African American voters, a population that could be consequential in local policy voting, especially given the large populations of African Americans in the areas where riots occurred.

Our findings may also differ because of the distinct context of the riots being studied—perhaps Los Angeles is unusually sympathetic to demands of the rioters. However, this explanation is unlikely given that Los Angeles County in the early 1990s was the type of place in which one would *a priori* expect a backlash against a riot. At that time, Southern California was considered a conservative stronghold in a Republican-leaning state (between 1900 and 1992, California elected only three Democrats as governor).

Another important consideration is that previous findings may be biased: despite making widespread claims about the effects of riots, most previous studies lack reliable evidence about whether riots actually cause the changes ascribed to them. To our knowledge, there have been no well-identified studies of the effects of violent protest on policy outcomes.[16] Rather, to understand the effects of riots, scholars have had to depend on post-treatment measures of single riots or cross-sectional studies of riots across multiple cities. Such cross-sectional studies must rely on the strong assumption that a riot's occurrence is uncorrelated with other variables that may affect politics downstream. Given that policy-oriented and scholarly accounts of riots treat them as non-random events caused by particular social and institutional forces (United States National Advisory Commission on Civil Disorders 1968; Olzak, Shanahan, and McEneaney 1996), this assumption seems implausible.

In fact, the use of cross-sectional data presents an acute threat of bias in the estimates of the effects of riots on behavior that may explain why previous studies have associated riots with a backlash. If there is a pre-existing trend in attitudes or policies to which rioters are

[16] One possible exception is Wasow (2016), who uses an instrument for riots following the death of Martin Luther King to measure the effects of riots on conservative voting.

reacting—that is, if riots are motivated by actions or policies that reflect locally unfavorable treatment or opinion by the majority, as the scholarly literature asserts— then local public opinion measured after a riot will be confounded by the conservative shift already underway.

3.8.2 Conclusion

Here we offer evidence that riots may change policy support locally in a direction sympathetic to the demands of the rioters. As such, we provide a nuanced understanding of how large-scale changes in American politics in the last half-century are linked to urban unrest, and how recent events may shape local politics in the future.

We have focused on violent protest as a political tool for a low-status group in the United States. While other scholarship has examined other forms of political action and asked if it is efficacious for racial minorities and other low-status groups (Verba et al. 1995), the scholarly literature has largely failed to ask whether rioting is a useful tool for accomplishing political goals, even though, from the perspective of the rioters, this question is paramount. Here we show that violent political protest can spur political participation among people who share an identity with the rioters.

Although it often seems extreme from the American perspective, political violence is not isolated to a particular time and place and is not rare in many parts of the world. Moreover, the implicit threat of violence underlies the relationship between governments and citizens in many places. As the use of violence continues to be an active feature of our political system, our findings and approach may help future scholars better understand this important topic.

4 Segregation drives racial voting: New evidence from South Africa (with Daniel de Kadt)

4.1 Introduction

From the United States to sub-Saharan Africa, groups remain physically separated along racial, religious, ethnic, or even economic lines. Segregation is sometimes the deliberate result of state policy (Davies 1967; Massey 1993); sometimes it is the result of "selection" into ethnic or racial blocs (Schelling 1971a; Telles 1992; Ihlanfeldt and Scafidi 2002; Zorlu and Latten 2009; Clark and Rivers 2013). Recent research considers the origins of segregation (Van Kempen and şule Özüekren 1998; Bayer, McMillan, and Rueben 2004; Bajari and Kahn 2005), and the effects of segregation and inter-group contact, be they psychological (Allport 1979), social (Oliver 1999; Kasara 2014; Bhavnani et al. 2014), economic (Massey, Condran, and Denton 1987; Boustan 2013), or political (Ananat and Washington 2009; Enos 2014). Other studies directly explore how local racial or ethnic context moderates political behavior (Gibson and Claassen 2010; Ichino and Nathan 2013; Kasara 2013).[1]

This study investigates how racial segregation, which shapes the racial context in which voters live, affects vote choice in post-Apartheid South Africa. We focus on the effect of sustained white isolation – which measures the probability of encountering a non-white (Massey and Denton 1988) – on white racial voting.[2] South Africa's elections have been marked by

[1]Replication data and materials are available on the Harvard Dataverse at Sands (2017d).

[2]The decision to focus on whites' voting behavior and survey responses is not free of normative concerns. We take this decision not because we believe white political behavior is any more important than non-white political behavior (this is trivially not true), but because black South Africans were highly mobile after the end of Apartheid. Given their position of political privilege and power, whites were confronted in the

racial voting since the advent of democracy in 1994; white South Africans typically vote for "white" parties such as the Democratic Party, the New National Party, and the recently conglomerated Democratic Alliance. In fact, scholars of the region often refer to South African elections as a "racial census" (Mattes 1995; Johnson and Schlemmer 1996; Lodge 1999; Ferree 2006, 2010). We find that racial voting among white South Africans increases as local white isolation increases. While we focus on ANC voting, it should be noted that all of our quantitative estimates are largely unchanged by using more expansive definitions of the dependent variable – "black party vote share" (ANC + Inkatha Freedom Party + Economic Freedom Fighters) or "white party vote share" (DA + Freedom Front).

We use three different empirical strategies, each using novel data. First, similar to studies that leverage the built environment to study segregation (Ananat 2011; Kasara 2014), we use natural physical geography as an instrument for sustained racial isolation – hills, valleys, and ridges that act as buffers and breaks preventing racial mixing. Using two different measures of "hilliness" we instrument for levels of contemporary white isolation at the political ward level ($n \approx 2,900$), conditional on demographics prior to the end of Apartheid.[3] Our second empirical approach utilizes new voting district data for the 1999 national election. We implement a ward-level difference-in-differences analysis to isolate the marginal effect of changes in the degree of racial isolation within political wards, while explicitly controlling for changes in the demographic levels of each ward. Finally, to address ecological inference concerns, we use approximately 39,000 geo-referenced survey observations over an 8-year period to test whether local racial isolation affects individual self-reports of voting intentions.

Across all three approaches the findings are consistent: Whites living among other whites are more likely to vote racially. At the ward level, greater white isolation in post-Apartheid South Africa appears to cause greater white racial voting – whites voting *against*

---

post-Apartheid period with changes to the demographics of their lived contexts, while black South Africans quite literally left behind old lived contexts and built new ones.

[3]For 2011 there were 4,277 wards in South Africa. We have balanced panel data for roughly 2,900 of these wards. Those we are missing are typically in the former homeland areas of South Africa, which by design have few whites.

the majority-black African National Congress. At the individual level, whites living in areas of high white isolation, where inter-racial mixing is limited, are *less* likely to vote across racial lines than whites who live in less segregated areas. Using 20 years of household sales data for the Gauteng region, as well as survey questions about racial sentiment, we are able to show that these results are likely not driven by residential sorting. We also present evidence that our results are not driven by a weak instrument.

Our study extends previous work on context effects and political behavior. We provide a rigorous empirical evaluation of race context effects. We triangulate across three different empirical approaches to tease out the effects of white isolation on white racial voting, and do so at *both* the aggregate level (using voting and census data), and the individual level (using survey data). We introduce new data that may be leveraged by other social scientists. Empirically rigorous research into the political legacies of exclusionary economic and geographic structures, of which Apartheid is a central example, is scarce. Our study begins to unpack some of these consequences, and provides insights and space for future research.

4.2 Theories of Racial Voting

While our study is, to our knowledge, the first to systematically link the profound demographic and geographic consequences of Apartheid to post-Apartheid political behavior, much high quality research investigates voting behavior in South Africa. This is particularly true of "racial voting", which has unsurprisingly received close attention. As articulated in Ferree (2006), past explanations typically focus on three central threads – "expressive identity voting", "racial heuristics", and "policy and performance". This trio of explanations provides a particularly insightful frame for our study, in that each suggests multiple ways in which context effects may affect racial voting.

4.2.1 Expressive Identity Voting

Expressive identity voting is commonly invoked as an explanation for political behavior in South Africa (Johnson and Schlemmer 1996; Louw 2000; Friedman 2004). Such a proposition is not unique – many researchers in various settings have argued that voters tend to prefer co-ethnic (co-racial) candidates (Young 1976; Kinder and Sears 1981; Horowitz 1985; Adida et al. 2016), and vote as ethnic blocs (Ishiyama 2012). One argument particularly relevant to our study is situational identity theory, which suggests that identities (ethnic, racial, or other) are multi-faceted, malleable, and subject to contextual activation (Young 1976; Bates 1983; Horowitz 1985; Posner 2004; Eifert, Miguel, and Posner 2010). Given the prevalence of such theories and the rich vein of evidence in their favor, it is plausible that local isolation or segregation, which shapes local context, may directly affect the *salience* of particular identities. This may then influence voting behavior, specifically voting behavior linked to identity.

Context may also affect inter-group trust, in turn influencing the prevalence of expressive identity voting (Kasara 2013). Such arguments draw on Allport's (1979) early work on "contact" and "threat". Allport (1979) argues that repeated "contact" with a new group may improve an individual's disposition toward members of the new group. On the other hand, a sense of "threat" might emerge with exposure to the new group, translating into a negative disposition toward its members. Living in close proximity to individuals of other races may shape how friendly, or threatened, one feels. It is plausible that re-calibrated racial sentiments influence behaviors, particularly when politics is already racialized. This explanation is appealing, and Ananat and Washington (2009), Hopkins (2010), Enos (2014), and Enos (2016) find evidence in the U.S. that context may affect opinions and turnout. It is arguable, however, that theories of "contact" and "threat" are both under-theorized and somewhat overfit to the U.S. case. That is, they appear to capture the effects of increased integration for the white *majority* only. How to generalize beyond the U.S. is not obvious.

4.2.2 Racial and Ethnic Heuristics

While it is hard to dispute the empirical regularity of identity voting, explaining its origins remains a field of active research. In the African context, the complexities of identity voting are born out in a series of ongoing debates about when and why ethnic or racial voting occurs (Posner 2004; Dunning and Harrison 2010; Eifert, Miguel, and Posner 2010; Bratton, Bhavnani, and Chen 2012; Ichino and Nathan 2013). From these debates emerges Ferree's (2006) second central thread – that South African voters make use of race as an heuristic, by which to infer who will likely benefit from the policies implemented by a party (Ferree 2006, 2010).

Similar to this take on race in South Africa, recent work argues that ethnic voting in sub-Saharan Africa is in fact an expression of economic interests, or heuristics about group interests, rather than any pure co-ethnic effect, and may be moderated by other social features (Lieberman and McClendon 2013; Carlson 2015). Such theories of racial or ethnic heuristic voting provide insights for those studying context effects. For instance, Ichino and Nathan (2013) argue, in probably the most influential recent work on context effects in African politics, that in Ghana, local context affects voting behavior through a rational self-interest mechanism. If voting is partly informed by expectations over targeted spending, ethnic (or racial) context should provide information about how politicians will act once in power. If a member of one ethnic group lives in an area surrounded by those from other groups, they may be more disposed to cross-ethnic voting, believing that this strategy improves their chance of receiving targeted spending.

4.2.3 Policy Voting

The third theme, policy and performance, suggests that voters regularly make their decisions on the basis of evaluations of party positions and government achievements (Mattes 1995; Mattes and Piombo 2001; Bratton and Mattes 2003). Outside of the South African case, researchers in sub-Saharan Africa have similarly argued that economic concerns can sometimes be more important than ethnic or identity concerns (Posner and Simon 2002; Weghorst and Lindberg 2013). Given this, the empirical regularity of observed racial or ethnic voting may be merely an artefact of economic voting. It may be the case that ideological or policy concerns correlate strongly with race, creating an observational equivalence between the two (Abrajano and Alvarez 2005; Abrajano, Alvarez, and Nagler 2008). In the South African case, for example, whites tend to be wealthier, and the "white" parties tend to offer platforms that focus on lower taxes, pro-business economic reforms, and diminishing the protection and power of labor (Ferree 2006). As such, racial or ethnic voting may simply result from the long term economic disparities that fall clearly along racial lines in South Africa.

If citizens' "racial voting" is a reflection of evaluations and judgements of policy and performance, racial context might affect racial voting. By changing perceptions of the economic reality around them, or changing the types of information voters access, racial context may influence vote choice.

4.3 Race and Space in South Africa

The previous discussion suggests not only that segregation and context effects are of global interest, but that the South African context is particularly appropriate for, and suited to study. South Africa has a well known history of both formal and informal segregation (Davies 1967, 1981; Sparks 1990; Christopher 1990, 1994, 2001; Worden 2011). From the early days of settler colonialism physical space was consciously divided up on the grounds of race and ethnicity. By 1913 the unified South African government was engaged in systematic legislative displacement of black South Africans, as well as legal segregation. That very year, the Natives Land Act cordoned off the majority of land in South Africa for whites, restricting property ownership and transfer rights for black South Africans.

In 1948 Apartheid became policy, and the country's physical space was further partitioned into different areas for different groups. Whites typically dwelt in racially isolated and homogenous communities, wealthier and better resourced than the rest of the country. Millions of black South Africans, roughly 70% of the population at the time, were physically displaced, dispossessed, and forced to live in cramped, densely populated, poorly resourced areas. The Group Areas Act of 1966 became the backbone legislation of broader Apartheid, aggressively regulating the right to live in certain areas, ownership and transfer of property, and even travel and commuting (Davies 1981).

In Johannesburg, the most populous city in South Africa, black South Africans lived in quasi-ghettos like Alexandra and the South-West Townships (Soweto), while whites remained in wealthy northern neighborhoods like Parkhurst, Houghton, or Sandton. These racial patterns were acts of policy; the Apartheid regime planned cities extensively, using physical barriers and buffer zones, insulating the white population from the larger non-white population (Davies 1981; Christopher 2001). Amenities, schools, and public facilities were also segregated (Christopher 1994). While small numbers of black, coloured, and Indian South Africans lived in white-designated areas, racial mixing on any larger scale was rare and short-lived, and shaped by white supremacist socio-political discourses that privileged whites over all others.

As the social and economic architecture of Apartheid collapsed in the early 1990s, pressure mounted for a negotiated transition to inclusive democracy. In 1990 F.W. de Klerk released Nelson Mandela from Victor Verster prison, and lifted the ban on the African National Congress (ANC). Parliament undertook a series of legislative reforms, dismantling legal Apartheid, ending formal segregation. In June 1991 the Group Areas Act was repealed, and the free movement of non-whites throughout South Africa was legally secured.[4] This temporal shift allows us to examine variation in the post-Apartheid period, but always conditional on Apartheid-era demographic profiles.

Black South African communities faced collective choices about post-Apartheid movement and settlement, constrained by the pre-existing arrangement of racial exclusion and

[4] The census data from 1991 was collected in May 1991, before the legislation dismantling the Group Areas Act and associated Acts was passed. Of course, there was some mobility in the cities before 1991.

ownership. In some cases they settled in lower-income high-rise areas like Yeoville and Hillbrow in Johannesburg. In others informal settlements sprang up in the middle of suburbs or cities. Cato Crest, which lies between the white areas of Glenwood and Westville in Durban is a prime example. By 1996, when the first post-Apartheid census was conducted, levels of segregation had declined greatly since 1991, but remained marked and obvious nonetheless (Christopher 2001).

Post-Apartheid mobility was in part determined by zoning regulations from the 1950s. These created "buffer strips" of empty space and used natural topographical and physical features (e.g. hills and rivers) to separate residential settlements of different race groups (Kuper, Watts, and Davies 1958; Davies 1981; Evans 1997; Christopher 2001). By law, buffers prevented contact between residential zones were a minimum of 200-500 yards, depending on the presence of roads and highways.[5] They were often much larger, and routinely created by displacing pre-existing communities and leaving land empty. Buffers remained empty and undeveloped during Apartheid, but were later settled after 1991 (Evans 1997).

Figure 4.1 shows the distribution of race by ward from the 2011 census data for Johannesburg and surrounds. The chloropleth shading indicates values on one measure of hilliness – *LogMeanTRI* – from low (light) to high (dark). The dots each represent 3,000 people, with blue dots representing black South Africans, and the red dots whites. The map visualizes the relationship described above. The darker hills and ridges running through the center of Johannesburg serve to segregate less densely populated, majority white areas, from more densely populated black areas. The mass of black population toward the bottom left is Soweto, the mass in the middle right Alexandra, and the mass toward the middle the topographically flat city center (and areas like Hillbrow and Yeoville). Re-integration patterns in post-Apartheid Johannesburg follow natural geography.

Topographic features partially determined the space on which to build shanty towns and informal settlements. They continue to serve their Apartheid-era purpose: circumscribing, to some degree, where black South Africans settle. The Durban settlement of Cato Crest

[5] "Guidelines for the Planning of Native Urban Areas" (NTS 4271 6 120/313)

Figure 4.1: Johannesburg by race (2011) and hilliness

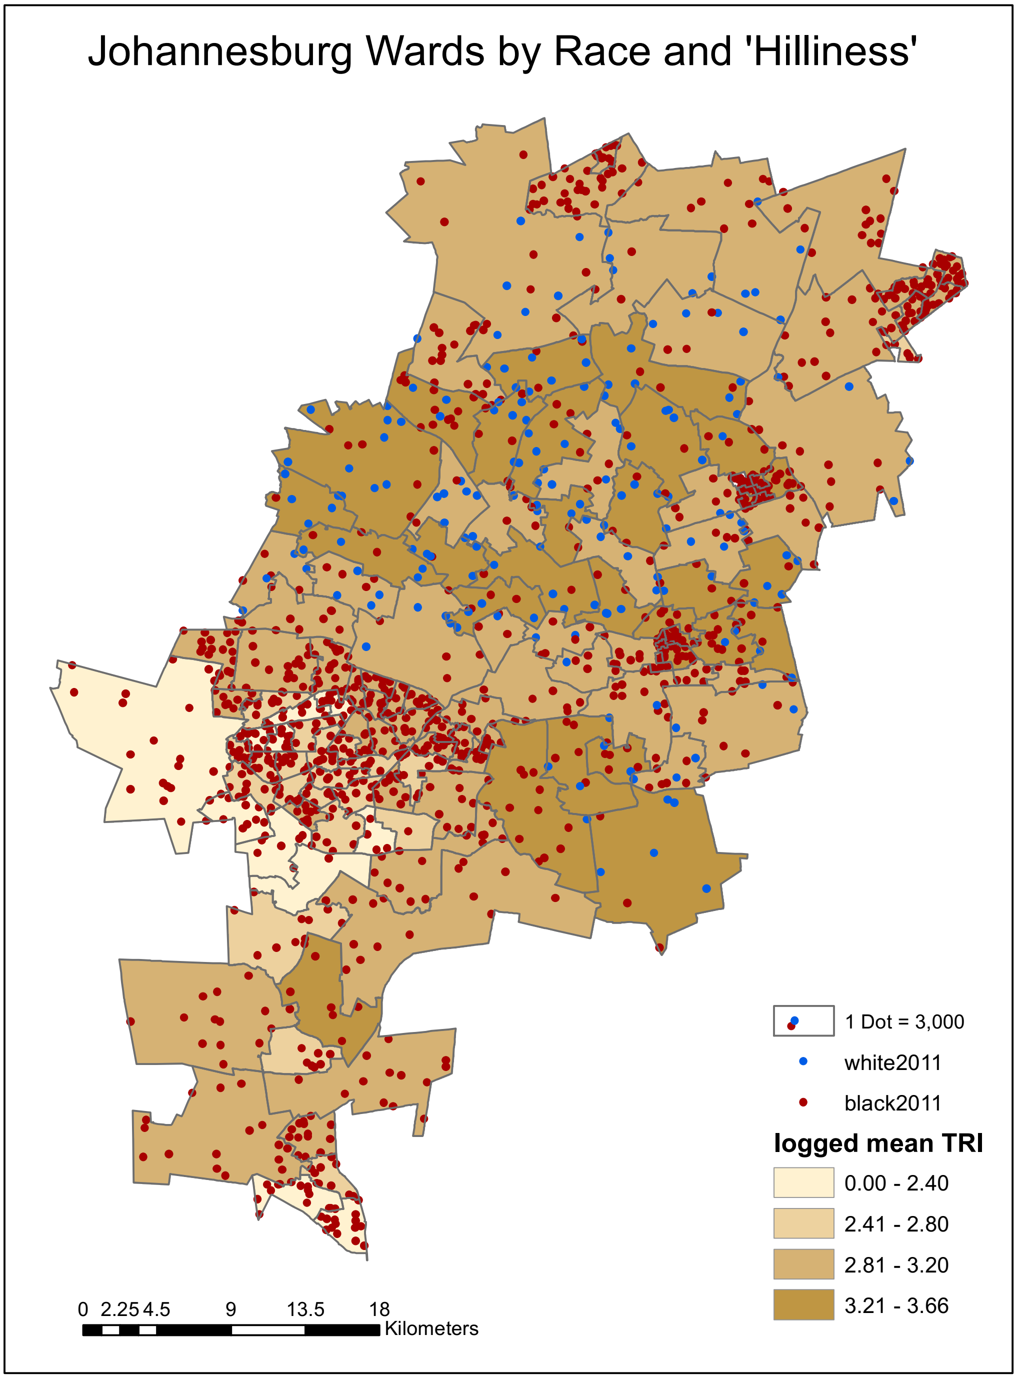

Note: This map depicts the metropolitan city of Johannesburg. The chloropleth shading indicates geographic hilliness, values of $LogMeanTRI$, from low (light) to high (dark). Each dot represents 3,000 people. Blue dots are black citizens, and the red dots white.

illustrates the point well. Cleared in the 1960s to create a vacant buffer zone (Edwards 1994), Cato Crest was a prime location for resettlement in the post-Apartheid period (Khan and Maharaj 1998). The area soon became oversubscribed with informal housing, shack-dwelling, and limited social services. Yet the physical space barely expanded – it is squeezed between two large physical ridges, the main Durban ridge and the hills of Westville, both predominantly white areas. These hillier neighborhoods were insulated from demographic change in part by the very physical geography they inherited from Apartheid. Areas like Manor Gardens and Glenridge, which lie in the flatter lands below the ridges and directly abut Cato Crest, integrated more rapidly. They became less segregated, despite, during Apartheid, being demographically and economically comparable to places that remained segregated.

Since the end of Apartheid the racial map has changed dramatically, but the degree to which Apartheid-era racial distributions persist today is a function of both geography and baseline (Apartheid-era) segregation. Our research design exploits this – we explicitly control for Apartheid-era demographics, and isolate variation in post-Apartheid racial distributions induced by topographic features. We can then study how this exogenous variation in segregation affects racial voting.

Contrasting Figure 4.1 with Figure 4.2 illustrates this intuition. Figure 4.2 shows the naive spatial association between white isolation (measured in 2011) and voting behavior (2014) in Johannesburg. The shading represents values of *WhiteIsolation* from low (light) to high (dark). The dots in each ward are sized proportional to the share of the 2014 vote that went to the ANC in that ward. Dark shaded regions, where the average white individual is surrounded primarily by other whites, support the ANC at lower levels than in lightly shaded regions, where the ANC garners most of the vote. Areas with high hilliness have correspondingly high levels of post-Apartheid white isolation, and also lower levels of voting for the ANC.

Figure 4.2: Johannesburg by ANC vote share (2014) and white isolation (2011)

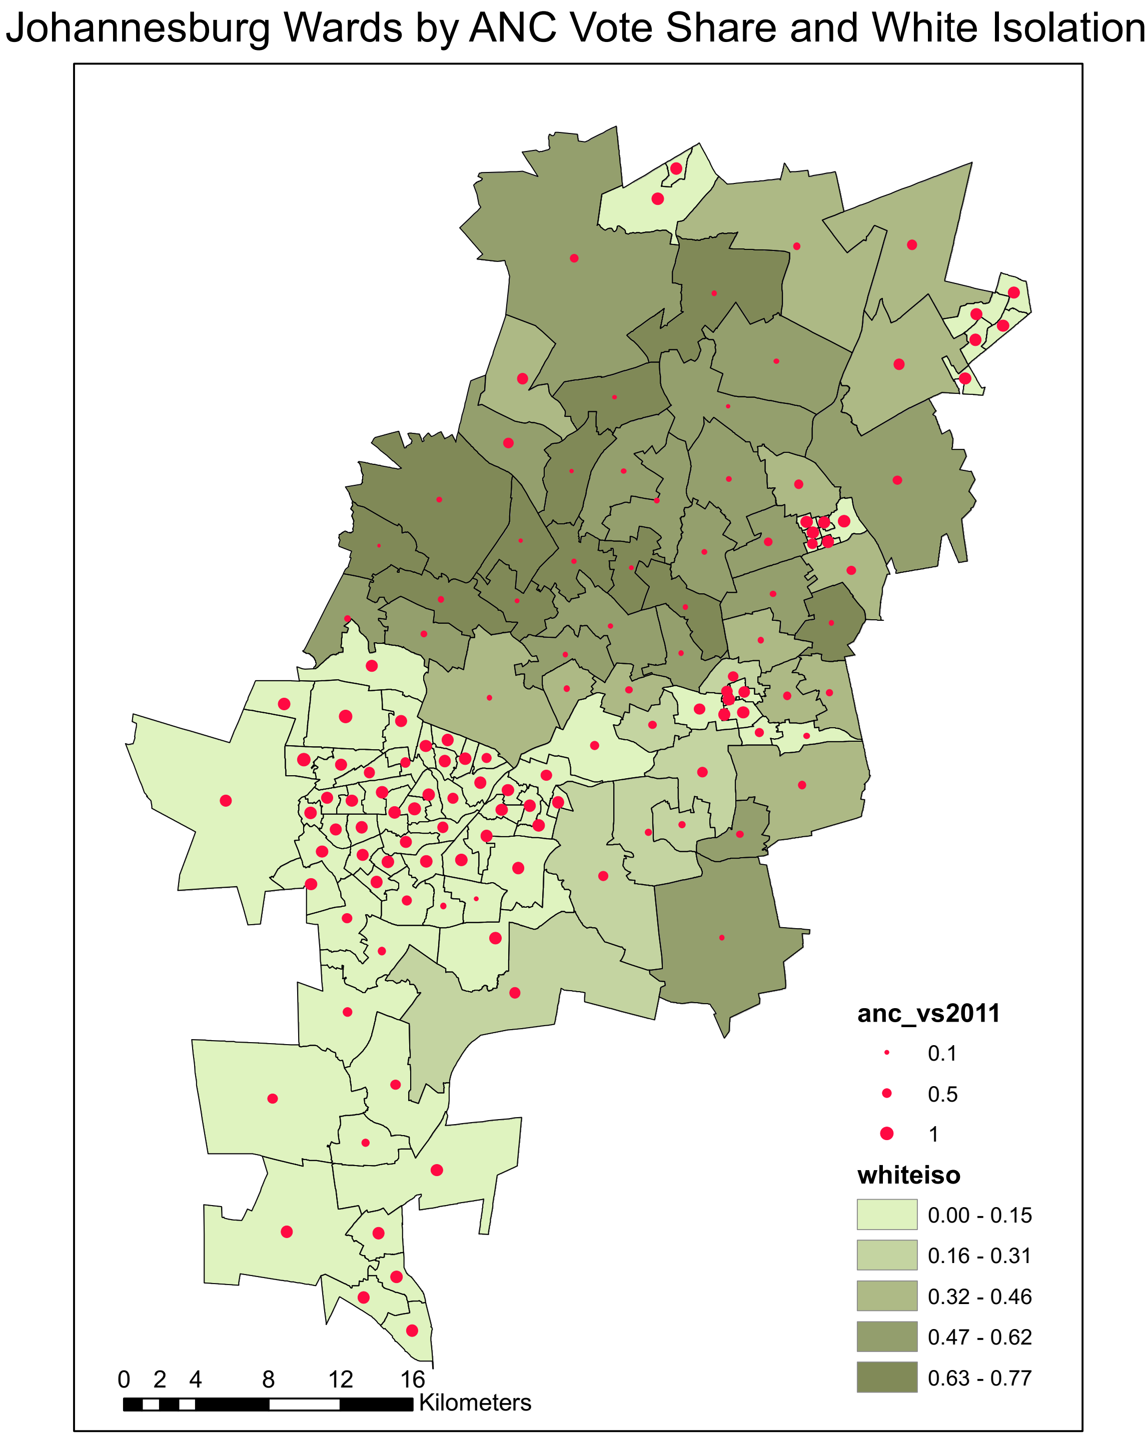

Note: This map depicts Johannesburg. The chloropleth shading indicates values on *WhiteIsolation* in 2011, from low (light) to high (dark). The red dots located in the centroid of each ward are proportional in size to ANC vote share in the ward in 2014. The largest dots show wards where the ANC attained nearly all of the vote and the smallest dots wards in which the ANC attained vote shares.

4.4 Empirical Strategy and Data

Voting in South Africa is racialized (Mattes 1995; Ferree 2006, 2010). While researchers find that black South Africans vote in less "racial identity"-centric ways than might be expected given the country's past (Mattes and Piombo 2001), white voters vote with high regularity for the Democratic Alliance (DA), the ostensibly "white" party. In our survey data around 75% of whites who do vote report a preference for the DA. The DA and the incumbent ANC soak up the majority of votes nationwide; in 2009 and 2014, they garnered over 80% of available votes. In the most recent local election, the vast majority of local municipalities were won by either party.

Given this political landscape, we measure racial voting among white South Africans as their propensity to vote *against* the ANC. While there are alternative non-white parties like the Inkatha Freedom Party (IFP), the Economic Freedom Fighters (EFF), or the United Democratic Movement (UDM), these parties have vanishingly small numbers of white supporters. They are, unlike the ANC, essentially black-only parties. As such they can be largely ignored in our analyses. Nonetheless, all of our quantitative estimates are largely unchanged by using as the dependent variable "black party vote share" (ANC + Inkatha Freedom Party + Economic Freedom Fighters) or "white party vote share" (DA + Freedom Front). Given this operationalization, if racial voting increases as local white isolation increases, whites living in areas in which they are more isolated are predicted to vote *against* the ANC at higher rates than those who live in mixed settings.

To test for context effects on white racial voting we constructed a new and unique panel-dataset at the lowest political level in South Africa, the ward.[6] At the aggregate level, we estimate the relationship between racial segregation and voting behavior *for wards*. We measure hilliness, segregation, and political behavior at the ward level, and make statistical comparisons across varying wards or over time. To do so we first use an instrumental

[6] In 2011 there were 4,277 wards in South Africa, electorally applicable for the 2011 local government and 2014 national elections.

variables approach, in which we instrument for current levels of segregation using hilliness, while controlling explicitly for Apartheid-era demographic profiles. The strategy is similar, though the unit of analysis more fine-grained, to Ananat's (2011) study of segregation in U.S. cities. While such an approach encounters classical ecological inference problems (for instance the modifiable areal unit problem or MAUP), the ward is the most high resolution aggregate unit of analysis feasible – at any lower unit of analysis either demographic or electoral data suffers severe measurement issues.

For each ward, we combine three types of data. First, we source census data from Statistics South Africa (StatsSA), the country's census bureau. We complemented the publicly accessible census data (available for 1996, 2001, and 2011), with a new dataset of measures from the 1991 Apartheid-era census, gathered shortly before the end of legislative Apartheid in mid-1991. Second, we source voting district electoral data for the National General Elections of 1999 and 2014, from the Independent Electoral Commission (IEC) of South Africa. Third, we source data on the physical geography of South Africa in a 1 km x 1 km grid, or raster layer. All data is then aggregated to the 2011 ward level, our fixed unit of analysis.

We give specific focus to data from two years – 1991, before the end of Apartheid, and 2011 (matched up with the most recent 2014 electoral data), the year of the most recent census. With this data we explore how, conditional on baseline Apartheid-era demographic profiles, post-Apartheid segregation or re-integration relates to voting behavior. Our panel dataset includes roughly 2,900 unique ward-level observations over two periods (1991 and 2011/2014), within roughly 217 municipal clusters.

4.4.1 Demographic data

To measure demographic features of each ward we rely on official census data. Our research strategy relies on having data from before the end of Apartheid, allowing us to condition on baseline levels of racial composition before the end of formal segregation. Due to the political and bureaucratic realities of the 1994 transition, historical census data has to this point been unavailable in digital form, as has geographic data linking old censuses to physical space.

In 2014 StatsSA made available the 100% sample of various historical Apartheid censuses, conducted from 1960 to 1991. We acquired a digital map for most of the 1991 enumeration areas (there are no maps for any earlier censuses at a high resolution level), and spatially joined this geography to the 2011 political ward geography, allowing us to locate each 1991 enumeration area (EA) within the larger ward boundaries. We linked each observation in the 100% sample of the 1991 census to its correct EA in 1991, cross-checking our matches manually for accuracy. We are able to match the 1991 individual-level census data with the 2011 wards, aggregating the 1991 data from the individual level up to the ward level. This gives us a range of 1991 census measures for the 2011 wards. The resulting dataset is unique in the study of South Africa. Variables and definitions of variables are available in the appendix.

Limitations remain. In the 1970s the government created independent Bantustans, four of which were considered "separate countries" by the Government. Censuses conducted before the end of Apartheid excluded the four independent Bantustans, mostly located in poorer rural parts of the country. The 1991 data thus refers to roughly 2,900 of the 4,277 modern wards.[7]

Given this data, we use *WhiteIsolation* as our central measure of ward-level "segregation". This is an appropriate measure in the South African context as it reflects the spatial *distribution* of race groups as well as the *proportion* of each group in a ward. Typically, segregation is thought of as the absence of "mixing" within a geographic space. For instance, within a single ward, there may be a 50-50 split in race, but complete racial separation within that space. Our measures take into account asymmetries in exposure and composition.

The *WhiteIsolation* variable was adapted from Massey and Denton (1988), and captures segregation *within* wards. To create this measure, we used census data at the EA level for the

[7] There were separate censuses for three of the Bantustans, but matching these with the available shapefiles has proven challenging. Fortunately, we feel confident that the missing observations are for the most part not relevant to this study. We include in every specification a municipality ($n = 234$) fixed effect, and so consider only variation *within* municipalities; as a result, rural municipalities in which there is no variation in levels of isolation, like the Bantustans, are of little inferential value.

years 1991 and 2011.[8] In any given census year, South Africa is divided into roughly 80,000 EAs, each containing approximately 650 people. To create the measure *WhiteIsolation*, for each EA, the white population residing in that area is divided by the white population in the ward and then multiplied by the white proportion of the total EA population. This value for each EA is then summed for each ward. More formally, it is equal to $\sum \frac{white_{EA}}{white_{ward}} \cdot \frac{white_{EA}}{total_{EA}}$. For example, if *WhiteIsolation* equals 0.42, this can be interpreted to mean that 42 of every 100 people a white person encounters will be white. The minimum value of *WhiteIsolation* is zero, where smaller values mean that the average white person resides in an area with a higher proportion of blacks. The maximum value of *WhiteIsolation* is 1, meaning that whites live in exclusively white areas.

4.4.2 Election data

We integrate into our dataset the election returns for the National General Election held in May 2014, again aggregated to the 2011 ward-level, the applicable wards for the period 2011-2015. Unfortunately we do not have election data at the baseline, for the 1994 election, but we are able to construct (for the first time, to our knowledge) ward-level returns for the 1999 National General Election.[9] While 1994 was a breakthrough democratic moment for the country, the election itself was plagued by fraud, miscounting, general administrative errors, and even alleged computer hacking by right wing groups. The South African government has never made the disaggregated data available to researchers or the public, possibly to maintain the credibility of the first election. We use the election data availble to generate ANC vote share variables, descriptions of which are available in the appendix.

[8] We aggregated the 1991 data available at the individual level into EAs. For 2011 the data is available at the Small Area Level (SAL). The SAL is, for the most part, exactly equivalent to the EA, with the exception of around 5% of the most sparsely populated EAs, which are sometimes aggregated up to the SAL to protect privacy. Throughout this paper we refer to the data as measured at the EA level for ease, but this proviso should be noted. We do not believe it has any ramifications for the results we present.

[9] The 1999 data is available at the voting district level ($n = 14,659$). We converted a shapefile of the voting districts to centroid points, then spatially joined this with the 2011 ward boundaries, and aggregated the voting data up to the 2011 ward level.

4.4.3 Geographic data

The degree of post-Apartheid segregation in any given ward in 2011 may be endogenous to voting behavior, even conditional on baseline covariates. To parse out exogenous variation in segregation in the post-Apartheid period, we leverage hilliness as an instrument for segregation. We calculate two different instruments to capture hilliness, and use both throughout the remainder of the paper. We include both variables to be transparent: measuring hilliness is difficult and largely unprincipled, so there is no one obviously correct measure – a number of measures could be justified. We also provide short discussions of the benefits and downsides of each measure.

Both measures are calculated from underlying data from the STRM30 Global Elevation Dataset, comprising CGIAR-STRM (3 seconds resolution) data aggregated up to 30 seconds. In practice, 30 arc-seconds (0.0083 degrees) corresponds to roughly 1 km^2 per pixel, or cell, in the raster layer. Thus our dataset contains an elevation value (or altitude) for every square kilometer in South Africa.

The first measure comes from a Terrain Ruggedness Index (*TRI*) which we calculate for each cell in the raster layer. The *TRI* is defined as the mean of the absolute differences between the value of a cell and the value of its eight surrounding cells.[10] The *TRI* was developed by Riley, DeGloria, and Elliot (1999) to quantify topographic heterogeneity in wildlife habitats, and has been applied elsewhere by social scientists as a measure of terrain ruggedness (Burchfield et al. 2006; Nunn and Puga 2012). The STRM30 data, which gives a unique altitude value for each square kilometer, can thus be used to generate, for each square kilometer of South Africa's landscape, a measure that captures wider heterogeneity of neighboring elevations.

To calculate the TRI, let $e_{r,c}$ denote elevation at point (r, c) in a grid of elevation points. The *TRI* at that point is equal to $[\sum_{i=r-1}^{r+1}\sum_{i=c-1}^{c+1}(e_{i,j} - e_{r,c})^2]^{\frac{1}{2}}$. Having generated *TRI*

[10] We also computed *Roughness*, developed by Wilson et al (2007), and defined as the difference between the maximum and the minimum value of a cell and its eight surrounding cells. This measure is very highly correlated with *TRI*, and the results are almost indistinguishable.

at the 1 square kilometer cell level, we then overlayed the ward boundaries on the raster layer, and aggregated up from the raster layer to the ward level. We take the mean of *TRI* for each ward, which is distributed log-normal, so we employ a logged transformation of the mean + 1 (to preserve positive only values).

1. *LogMeanTRI*: This is the central tendency of the *TRI* measure within a ward. A more granular measure of hilliness, it represents the average degree of variety in a ward's terrain. However, because *TRI* is best suited to capture fine-grained variation, the mean value in the ward fails to describe the spread of terrain across a large and varied area, and is sensitive to outliers.

The second measure differs from the first by incorporating features of the surrounding wards, rather than focusing on the given ward itself. This serves to capture the idea that it may be that physical boundaries around a ward have consequences for segregation within that ward:

2. *CNLRA*: This is the mean of a measure we term *LogRangeAlt* for the *contiguous neighbors* of any given ward. *LogRangeAlt* is the difference between the highest cell in a ward and the lowest cell. A high value means that, somewhere in the ward, there is a very high point, and a very low point. A low value means that all the cells in the ward are of similar altitude. This is an elegant measure of hilliness in small wards, where there are only a few square kilometer cells, but is weaker when the wards are larger, as it may overstate hilliness. Because we focus on contiguous neighbors, rather than measuring the hilliness of a given ward, *CNLRA* measures whether a given ward is *encircled* by other hilly wards.

4.5 Aggregate Analyses

Assessing the relationship between segregation and voting behavior is challenging; there may exist many potential confounders – omitted variables that jointly predict both segregation and voting behavior, yielding spurious estimates. To navigate this empirical hurdle, we use instrumental variables (IV) and difference-in-differences methodologies.

4.5.1 Instrumental Variables

The IV approach is our primary research design, exploiting the context outlined in section 4.2. We consider hilliness as an "encouragement" – higher values of hilliness encourage higher levels of isolation or segregation in the post-Apartheid period. Segregation, measured as white isolation, is our "treatment". We assess how this affects ANC vote share by "instrumenting" for isolation with hilliness; this allows us to view the treatment as the change in segregation within a ward that results from the confluence of the end of Apartheid and hilliness, which can be viewed as exogenous or as-if randomly assigned. This approach allows us to isolate the effect of segregation, eliminating numerous potential confounders.

The approach hinges on two central assumptions:

1. Conditional ignorability of hilliness: Levels of hilliness are conditionally uncorrelated with the potential outcomes of current levels of segregation, and potential outcomes of voting behavior.

Assumption (1) allows us to isolate the relationship between hilliness and contemporary levels of segregation, as well as the "intent-to-treat" effect of hilliness on voting behavior. To isolate the relationship between segregation and voting behavior, which is our key interest, we require a second assumption:

2. Exclusion restriction for hilliness as an instrument: This states that hilliness only affects political behavior by affecting current levels of segregation.

Assumption (2) requires that our instrument affects our outcome of interest through no channel other than the treatment, white isolation. In this case it requires that hilliness not have an effect on political behavior outside of the effect channelled through segregation. In Section 4.7 we discuss the validity of these two assumptions, and present a number of tests that suggest they are met.

To estimate the IV regression, we run two specifications. The first stage is:

$$WhiteIsolation_{i,2011} = Hills_i\beta + \mathbf{X_{i,1991}}\gamma + \mathbf{W_{i,1991}}\omega + \delta_m + \epsilon_m$$

Where $WhiteIsolation_{i,2011}$ represents the extent to which whites are separate or isolated from blacks, for a given ward i in 2011. $\mathbf{X_{i,1991}}$ is constant across both the first and second stages, and represents a set of baseline demographic covariates from the 1991 census, for a given ward. In particular, we include the baseline measure of segregation, as well as the proportions of whites, blacks, and coloureds, in the ward in 1991. In some specifications we also include $\mathbf{W_{i,1991}}$, a set of economic covariates from 1991 including education, income, and employment. To ensure that we focus only on variation *within geographic units* rather than across them, we include δ_m, which are municipality fixed effects for each municipality m (there are roughly 217 municipalities for the 2,900 wards in our data). The errors, ϵ_m, are clustered at the municipal level. We then introduce a second stage regression using the fitted values from the first stage. This two-stage least squares regression model allows us to investigate the political consequences of sustained segregation and suppressed diversity by partialing out the endogenous portion of $WhiteIsolation_{i,2011}$:

$$Y_{i,2011} = \widehat{WhiteIsolation}_{i,2011}\beta_2 + \mathbf{X_{i,1991}}\gamma_2 + \mathbf{W_{i,1991}}\omega_2 + \delta_m + \upsilon_m$$

Where Y represents political outcomes of interest (vote choice). $\widehat{WhiteIsolation}_{i,2011}$ is the fitted value of $WhiteIsolation_{i,2011}$ for a given i in 2011, which comes from the first stage regression. The first stage regression includes $Hills_i$, which is one of the three measures of hilliness outlined above, for ward i (note that hilliness is of course time-invariant). $\mathbf{X_{i,1991}}$ is once more a set of baseline demographic covariates from the 1991 census, for a given ward, while $\mathbf{W_{i,1991}}$ is an optional set of economic covariates. We again include δ_m, which are municipality fixed effects for each municipality m. As before, the errors, υ_m, are clustered at the municipal level.

β_2 is the quantity of interest, which is the partial effect of variation in $WhiteIsolation$

that is due to exogenous variation in $Hills$. Given that we condition on white and black proportions in 1991, the coefficient β_2 can be understood as the added effect of increased segregation on political outcomes.

The results of the first stage regressions can be found in the appendix. Both measures of hilliness induce variation in our measure of segregation, satisfying the relevance assumption of the instrumental variables identification strategy.

We present in Table 4.1 the results of our IV analysis. The dependent variable is ANC vote share in 2014, and the endogenous treatment is white isolation in 2011. The results suggest that racial isolation of whites has a strong, statistically significant effect on voting behavior. The effect is robust, shown in columns (3) and (4), to the inclusion of baseline economic covariates.

To interpret the magnitude of these effects, consider the coefficient -0.690 in the first row of column (3) in Table 4.1. This suggests that as the predicted level of segregation in the ward, proxied in the first stage by the impact of *LogMeanTRI* on the white isolation index, increases by 0.10, or 10%, ANC vote share decreases by 0.069, or 6.9 percentage points. This represents a decrease of a quarter of a standard deviation, a substantively and statistically significant effect. Moving from the first to the third quartile of white isolation is associated with a full standard deviation change in ANC vote share.[11]

4.5.2 Difference-in-differences

While the IV provides evidence that higher levels of white isolation induce voting against the ANC, one might argue that these effects are merely evidence *of* racial voting, rather than evidence that racial voting is affected by segregation. Our measure of white isolation correlates with the fraction of the population that is white, so an increase in the isolation index should lead to an decrease in ANC vote share. In the IV specifications we resisted including any covariates from after 1991 to avoid the possibility of post-treatment bias. A

[11] The difference between the first quartile of white isolation and the third quartile is 0.44 on a 0 to 1 scale. Thus we would predict ANC vote share to decrease by around 0.3, or 30%

Table 4.1: Instrumental Variables Regression of White Isolation (2011) on ANC Vote Share (2014)

	Dependent variable:			
	ANCvs2014			
Instrument	*LogMeanTRI*	*NN_LRA*	*LogMeanTRI*	*NN_LRA*
WhiteIsolation2011	-0.615	-0.366	-0.690	-0.469
	(0.291)	(0.180)	(0.346)	(0.245)
WhiteIsolation1991	0.173	0.0938	0.185	0.120
	(0.0899)	(0.0579)	(0.0980)	(0.0704)
WhiteFrac1991	0.00573	-0.0929	0.0498	-0.0472
	(0.107)	(0.0736)	(0.141)	(0.0985)
BlackFrac1991	0.351	0.337	0.347	0.334
	(0.0654)	(0.0712)	(0.0605)	(0.0645)
ColFrac1991	0.00279	-0.00722	-0.00232	-0.0125
	(0.0752)	(0.0784)	(0.0743)	(0.0749)
First Stage F	10.71	24.62	7.76	13.17
N	2,866	2,898	2,850	2,882
R^2	0.690	0.704	0.688	0.710
Racial $\mathbf{X_{i,1991}}$	✔	✔	✔	✔
Economic $\mathbf{X_{i,1991}}$			✔	✔
Municipal FE	✔	✔	✔	✔
Cluster robust standard errors in parentheses				

difference-in-differences methodology allows us to estimate the interaction between white isolation and the fraction of the population that is white. Ideally this approach requires baseline electoral data, which, as discussed earlier, we do not have. We use $ANCvs1999$, the earliest election data for post-Apartheid South Africa available that can be constructed at the ward level, as a proxy for the vote share in 1994.

Difference-in-differences allow us to test whether more segregated areas vote less for the ANC, conditional on changes in the fraction of the population that is white. It has the added benefit of eliminating time-invariant or slow-moving confounders at the ward level. We estimate this in two ways, seeking to explain change in ANC vote share at the ward level between 1999 and 2011. First, we take the change in white isolation from 1991 to 2011, and interact this with the change in white fraction. The quantity of interest is the coefficient on this interaction term. This estimates how the association between the white fraction variable and ANC vote share changes as white isolation changes.

Second, we created a dummy variable for whether the level of white isolation in a given year (1991 and 2011) was greater than the median level in 1991 (0.09). We then created a "treatment" variable coded as the change in these dummies, such that places which switched from below the 1991 median to above it in 2011 are coded 1, places which stayed the same are 0, and places which switched from above the 1991 median to below it in 2011 are coded -1. That is, places that became more isolated take the value of 1, while places that stayed the same take 0, and places that became less isolated take -1. While somewhat crude, this approach makes interpreting the results of the difference-in-differences much easier.

We estimate the quantity of interest for both approaches using the following first-differences regression specification:

$$\Delta ANCvs_i = \alpha \Delta WhiteFrac_i + \beta \Delta WhiteIsolation_i + \gamma \Delta WhiteIsolation_i \cdot \Delta WhiteFrac_i + \zeta \Delta X_i + \delta_m + \epsilon_m$$

Where i indexes a ward, Δ indicates a first-difference (from 2011 to 1991) for any covariate, and $WhiteIsolation_i$ is either the dummy or continuous variable. X_i are covariates, δ_m

municipal fixed effects, and ϵ_m errors clustered at the municipality level. Our key quantity of interest is γ, which gives the difference-in-differences – how the relationship between the proportion of whites in an area and ANC vote share varies by levels of white isolation (treatment status). Notwithstanding the limitations imposed by our proxy for 1994 ANC vote share, this approach eliminates time-invariant confounders, and isolates the effect of white isolation if there are parallel trends between treated and untreated units.

The results of this exercise are presented in Table 4.2, with the sixth and seventh rows of coefficients being the key difference-in-differences estimates. As in the IV case, the coefficients are negative and statistically significant, implying that an increase in isolation increases anti-ANC voting, even conditional on changes in the proportion of a ward that is white.

Consider the coefficient for γ presented in column (4). For a given change in the proportion of a ward that is white, moving from below the 1991 median in white isolation to above the median decreases the ANC's vote share by around 12 percentage points. Consistent with this result, in column (2), where the treatment is continuous, moving from the lowest value of white isolation to the highest value induces a decrease of around 16 percentage points.

Together the IV and difference-in-differences results paint a consistent picture. Conditional on baseline characteristics, higher levels of white isolation (or segregation) in the post-Apartheid period predict decreases in ANC vote share. The results are not merely evidence of racial voting per se, but evidence of *racial voting becoming more intense in more segregated areas.*

4.6 Survey Analyses

At the ward level, for a given sized white population, more segregated areas tend to "vote white" more readily. But do *individuals* residing in areas that are more or less segregated behave differently than their counterparts in more or less mixed areas? If segregation combined with the end of Apartheid generates persistence in ward-level voting patterns, we

Table 4.2: Difference-in-Differences Regression

	Dependent variable:			
	$\Delta ANCvs$			
	(1)	(2)	(3)	(4)
$\Delta\,WhiteFrac$	-0.0644	-0.0488	0.00685	0.0128
	(0.0518)	(0.0517)	(0.0420)	(0.0437)
$\Delta\,BlackFrac$	0.157	0.169	0.164	0.172
	(0.0399)	(0.0429)	(0.0399)	(0.0431)
$\Delta\,ColFrac$	0.140	0.151	0.148	0.156
	(0.0403)	(0.0422)	(0.0398)	(0.0421)
$\Delta\,WhiteIsolation$	-0.00260	-0.00205		
	(0.0191)	(0.0191)		
$\Delta\,WhiteIsolationDummy$			0.000748	0.00226
			(0.00808)	(0.00825)
$\Delta\,WhiteFrac \cdot \Delta\,WhiteIsolation$	-0.195	-0.163		
	(0.0755)	(0.0742)		
$\Delta\,WhiteFrac \cdot \Delta\,WhiteIsolationDummy$			-0.146	-0.118
			(0.0506)	(0.0471)
N	2,824	2,809	2,824	2,809
R^2	0.552	0.552	0.552	0.552
Racial $\Delta\mathbf{X}$	✔	✔	✔	✔
Economic $\Delta\mathbf{X}$		✔		✔
Municipal FE	✔	✔	✔	✔
Cluster robust standard errors in parentheses				

should expect to find evidence of those patterns among individual voters' survey responses.

To examine context effects at the individual level, we spatially linked survey data from the South Africa Social Attitudes Survey for 2004 - 2011 to the ward-level data described earlier. The resulting dataset contains over 39,000 individuals for 8 periods.[12] We estimate the relationship between segregation on vote choice as a linear probability model:

$$ANCvote_i = \alpha + White_i\beta_1 + WhiteIsolation_w\beta_2 + White_i * WhiteIsolation_w\beta_3 + \mathbf{X_{i,w}}\gamma + \delta_m + \delta_t + \epsilon_w$$

ANC_vote_i represents individual i's self-reported vote choice; a vote for the ANC equals 1, and 0 otherwise. $White_i$ is an indicator for whether the race of the respondent is white, and $WhiteIsolation_w$ is the value of the white isolation index in individual i's ward w in the year 2001. $\mathbf{X_{i,w}}$ is a vector of individual-level covariates, including race, age, age squared, sex, marital status, educational attainment and wealth, and a set of ward-level covariates, including racial composition pre- and post-Apartheid, population density, employment rate and average education. The ward-level baseline covariates are measured in 2011. We also include municipality fixed effects, δ_m, and survey year fixed effects, δ_t. The error term, ϵ_w, is clustered at the ward level.

We repeat the exercise for both prospective (rather than retrospective) vote choice, and turnout, to assess whether the aggregate results are driven by mobilization rather than changes in vote preferences. In columns (1) through (9) of Table 4.3, the primary coefficient of interest is estimated on the interaction term $White_i * WhiteIsolation_w$. This gives the differential impact of local white isolation for whites on their likelihood of voting for the ANC. In particular, it captures the marginal conditional association between ANC vote and the degree of isolation in a ward for white individuals who reside in that ward.

Columns (1) through (3) suggest that whites living in greater isolation are especially likely to report having voted against the ANC in the prior election. Although whites in even

[12] In columns (1) through (6) of Table 4.3, we subset to self-reported voters only, about 60% of respondents. When we control for 1991 covariates the number of wards in the dataset drops considerably as we lose areas with no data in 1991, leaving approximately 18,700 observations. In columns (7) through (9) we include non-vters, and attrition in the sample is due to missingness of covariates and 1991 data.

Table 4.3: Individual Interaction Results for ANC Voting

	Dependent variable:			Dependent variable:			Dependent variable:		
	Retrospective ANC Vote			*Future ANC Vote*			*Turnout*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$Coloured_i$	-0.404	-0.389	-0.346	-0.414	-0.401	-0.363	-0.122	-0.158	-0.135
	(0.0147)	(0.0146)	(0.0180)	(0.0153)	(0.0150)	(0.0187)	(0.0105)	(0.0102)	(0.0121)
$Indian_i$	-0.387	-0.361	-0.329	-0.429	-0.392	-0.368	-0.190	-0.270	-0.255
	(0.0198)	(0.0202)	(0.0216)	(0.0184)	(0.0188)	(0.0207)	(0.0129)	(0.0119)	(0.0157)
$White_i$	-0.802	-0.772	-0.734	-0.748	-0.699	-0.666	-0.119	-0.262	-0.234
	(0.0262)	(0.0263)	(0.0294)	(0.0243)	(0.0249)	(0.0291)	(0.0260)	(0.0252)	(0.0278)
$WhiteIsolation_w$	0.0143	0.0246	0.0416	-0.00228	0.00816	0.0118	-0.0176	-0.0266	0.0229
	(0.0137)	(0.0137)	(0.0259)	(0.0165)	(0.0169)	(0.0309)	(0.0131)	(0.0118)	(0.0200)
$White_i * WhiteIsolation_w$	-0.111	-0.125	-0.147	-0.0525	-0.0781	-0.110	0.0518	0.0725	0.0402
	(0.0426)	(0.0422)	(0.0474)	(0.0409)	(0.0409)	(0.0488)	(0.0444)	(0.0414)	(0.0455)
N	26,383	24,134	18,663	26,383	24,134	18,663	35,269	32,184	25,390
R^2	0.478	0.491	0.498	0.392	0.402	0.425	0.051	0.248	0.244
Indiv **X**		✔	✔		✔	✔		✔	✔
Ward **X**			✔			✔			✔
Municipal FE	✔	✔	✔	✔	✔	✔	✔	✔	✔

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure 4.3: Marginal effects for (retrospective) ANC vote

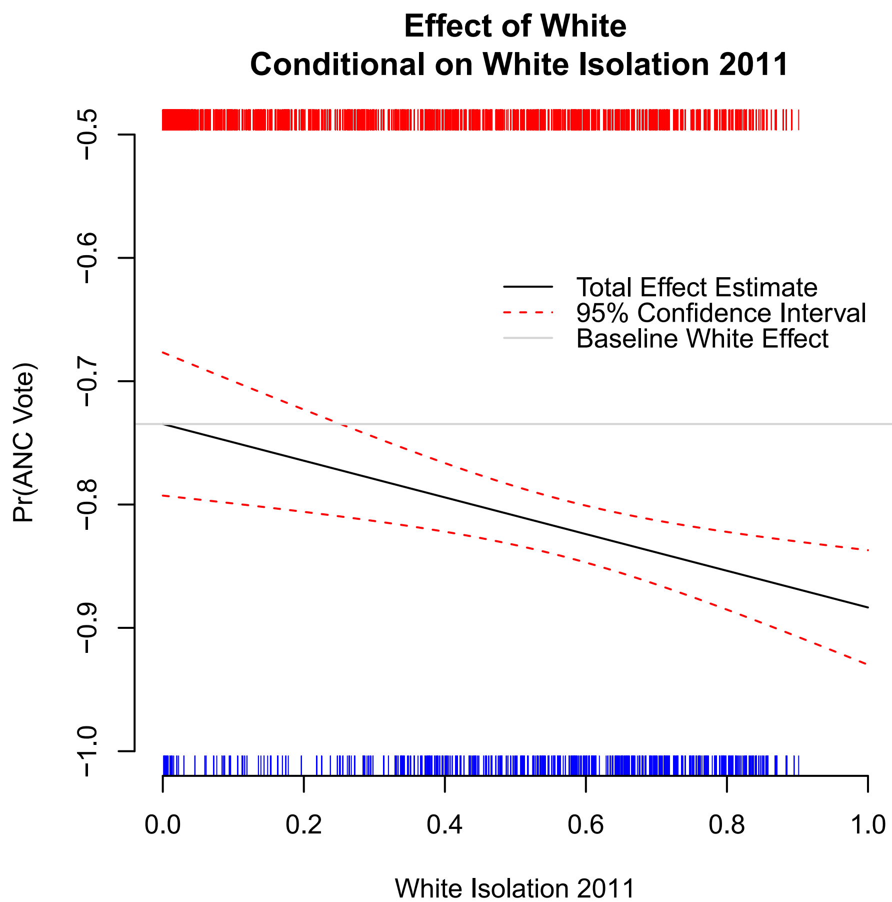

Note: This figure shows the predicted probability of voting for the ANC for white South Africans as a function of ward-level racial isolation (with black South Africans as the reference category). The curve and confidence intervals correspond to column (3) in Table 4.3. The lower (blue) rug-plot shows the distribution of white individuals with respect to $WhiteIsolation$. The upper (red) rug-plot shows the distribution of non-white individuals with respect to $WhiteIsolation$.

the least isolated wards are less likely to vote ANC compared to their non-white counterparts, this relationship is magnified among whites in whiter areas. For example, column (3) in Table 4.3 suggests that while whites in the first quartile of isolated wards are 73.6 percentage points less likely to vote for the ANC compared to members of other race groups, those in the third quartile of most segregated urban wards are an *additional* 7.1 percentage points less likely to vote for the ANC, controlling for a battery of individual-level and ward-level covariates.[13] The results of column (3) in Table 4.3 are visualized as a marginal effects plot in Figure 4.3. The figure shows the predicted probability of voting for the ANC for white South Africans as a function of ward-level racial isolation, with black South Africans as the omitted reference category.

Columns (4) through (6) reveal more modest estimates of context effects on prospective racial voting, but the pattern remains. Typically at statistically significant levels, living in a more white isolated context produces a higher degree of prospective racial voting for whites. Columns (7) through (9) repeat these tests for self-reported prior turnout, and while the coefficients are positive, they are less than half the magnitude of the retrospective vote choice results reported in columns (1) through (3), and typically not statistically significant. This suggests that the results cannot be explained purely by mobilization.

4.7 Discussion

Thus far we have presented evidence that white South Africans vote "more white" when they live in racially isolated communities. This result emerges across three different empirical strategies, and at two different levels of analysis. The IV approach, implemented at the ward level, eliminates numerous confounders by isolating "exogenous" variation in racial isolation in the post-Apartheid period. The difference-in-differences approach insulates our estimates from unit-specific time-invariant confounders, and allows us to explicitly control for the

[13]The first quartile of *WhiteIsolation* is 0.0036; the third quartile is 0.4778. Thus, all else equal, an individual at the first quartile is $(-0.735) + (0.0036 * -0.149) = -0.736$, or 73.6 percentage points less likely to support the ANC, and an individual at the third quartile is $(-0.735) + (0.4778 * -0.149) = -0.806$, or 80.6 percentage points less likely to support the ANC, compared to a non-white individual.

white fraction of the population in 2011. Finally, the individual level results suggest that our findings are not driven by an ecological fallacy.

What explains these remarkably consistent findings? First, we address the possibility that our findings are not causal at all, but instead the result of "sorting" – that certain types of people choose to live in certain types of places. We provide a range of evidence that strongly suggests sorting does not account for our findings. Second, using the theoretical frames introduced in section 4.2, we discuss various plausible explanations of the racial context effect in South Africa.

4.7.1 Ruling out sorting

The results may be threatened if, conditional on the demographic profile of an area in 1991, certain types of people *within each race group* were more likely to choose to live in more hilly, or more segregated, areas. Such sorting would violate the assumptions underlying our IV estimates, and could threaten our difference-in-differences and individual-level analyses.

One example of such a threat is the well-documented negative association between "rugged" terrain and economic outcomes (Nunn and Puga 2012). Irregular and sloped terrain poses challenges for cultivation and irrigation, and increases the costs of building and transportation. This may threaten our empirical strategy if economically less well off whites, also less likely to support the ANC, chose to settle in more hilly areas during Apartheid. These adverse features of ruggedness are more relevant in rural areas, where the success of agriculture is directly linked to land quality. In densely populated regions, it is less obvious that hilliness should be associated with wealth. As such, the primary risk comes from comparisons *between* rural and urban areas using our instrument. This concern is mitigated in our specifications by the inclusion of the municipal fixed effect, which ensures that comparisons are made within circumscribed geographic areas, either urban or rural.

Furthermore, if particular types of individuals choose to live in more or less segregated areas, associations should exist between hilliness and economic variables. Results are presented in the appendix that show very limited conditional associations between our measures of hilliness and covariates in 1991. There are few statistically significant associations, and where they do exist the magnitudes are small, increasing our confidence in the plausibility of the IV assumptions.

We also find stability in correlations between our instruments and these economic covariates over time. If the correlation between the instrument and demographic covariates shifts over time, we might suspect that hilliness affected political behavior through alternative channels. If the correlations remain stable, this is suggestive of limited change in observable covariates, and plausibly unobservables too. The results of this exercise are presented in the appendix. Comparing the same covariates in 1991 and 2011 shows minimal differences in the way our instruments correlate with economic covariates over time, none of which are statistically significant. This stability suggests hilliness is not predicting changes in economic covariates over time, providing additional evidence that the exclusion restriction is not violated.

We also probe sorting in the post-Apartheid period using new data on property deeds transfers in each enumeration area (EA) from 1993 to 2010 in Gauteng province and surrounds.[14] This dataset allows us to assess the extent to which South Africans move on the basis of race and/or hilliness when Apartheid came to an end. Fine-grained real estate transfer data is rarely available in developing contexts; thus this dataset is extraordinary in its level of detail and historical scope.

Although we do not have property transaction data for the entire country, the dataset includes at least a portion of five of the nine provinces, an area home to at least 5.5 million people in 1991 (around 20% of South Africa's population at the time), and includes Johannesburg, the country's most populous city. We link this dataset, an EA-level panel of the counts of every property within various price bands that changed hands in each year, to the 1991 census EAs as well as the ward-level data.

We first explored whether property sales in 1993 are correlated with the racial or socioe-

[14]This data was generously provided to us by Lightstone Property and the Gauteng City-Region Observatory.

conomic characteristics of neighborhoods (at the EA level). If large scale residential sorting occurred at the end of Apartheid, we would expect to see greater numbers of deed transfers, normalized by population, in areas with a higher fraction of whites. We find (results in the appendix) no evidence of "white flight" in 1993, nor do we observe a significant relationship between property sales and wealthier EAs or EAs with higher proportions of Afrikaans speakers. In fact, overall residential mobility in 1993 and subsequent years is typically very low; in 1993 only 5% of the existing formal dwellings in our coverage area were sold or transferred.[15]

We then aggregated the EA panel to the ward level in order to test whether property sales confound our main effect estimates. We regress segregation in 2011 on our measures of hilliness, deeds transfers, and the baseline covariates and show that the coefficients on our measures of hilliness are unchanged by the inclusion of property sales (results in appendix). We also regressed the number of deed transfers per person on each of our instruments and a set of baseline characteristics, including the 1991 fraction of each race group. In general, we do not find a significant conditional association between deeds transfers and our instruments (results in appendix). This finding is robust to the inclusion of various sets of covariates and municipality fixed effects; it is also robust to the set of years that we include in the dependent variable. It implies that none of our instruments predict property sales, a key measure of mobility and residential sorting.

Finally, we replicated our individual-level analyses reported in section 4.6, but using measures of racial sentiment as the dependent variable, rather than voting behavior. We focused on four questions, whether respondents felt "positively" toward whites (blacks) and whether respondents felt "friendly" toward whites (blacks). These analyses (results in appendix), reveal coefficients of approximately zero on the interaction terms for all four measures. That is, whites living in areas of higher white isolation are *not* more likely to report different levels of racial sentiment toward either black or white South Africans. One obvious refrain is that

[15]This figure is calculated by dividing the total number of deeds transfers in 1993 by the total number of formal households in the 1991 census, for the EAs covered by the deeds data. This is a conservative estimate given that we consider formal households only.

these measures are plagued by social desirability bias. While this is possible, we would note that there is a great deal of variation on all four measures, and that the base term point estimates for whites behave as would be expected – whites are more positive and friendly toward whites, and are less positive and friendly toward blacks. In sum, these various checks suggest sorting does not appear to have occurred on a large scale, and likely cannot account for our results.

4.7.2 Mechanisms

If our findings are not driven by residential sorting, what can explain the effects of racial segregation on racial voting in South Africa? By definition, segregation or racial isolation shapes the demographic make-up of local communities, directly affecting the diversity of individuals' lived contexts. This variation in context may bleed into behaviors, particularly in situations where race is already a salient or divisive issue. For instance, segregated local contexts may have particularly acute consequences for political behavior when parties separate along racial lines, as is the case in South Africa. As outlined in section 4.2, the three broad approaches to racial voting in South Africa provide insights into potential context effect mechanisms.

First, expressive identity voting posits that voting is an act of self expression; whites will vote for white parties if they see race as an important part of their identity, and make political decisions on the basis of this. This theory also lays the foundation for *situational* identity theories, which suggest that identities are subject to contextual activation. In the South African case, it seems plausible that white voters surrounded by members of the same race group may come to think of themselves as "more white" than those who live in more diverse contexts. Given the history of white supremacy and oppression in the country, race groups in South Africa are, in Horowitz's (1985) words "ranked" – in such a setting the powerful group may self-identify more strongly if they maintain separation and isolation. In turn, whites whose racial identity is highly salient may vote "more white". Such a sequence would explain our findings – voter context shapes the salience of voter identities, which determine voting behavior. For whites residing in more mixed areas, repeated contact with black South Africans may diminish the salience of their white identity. While this theory certainly has appeal in the South African case, our reading of the literature is that expressive identity voting in South Africa has largely been "debunked" in recent scholarship (Ferree 2006, 2010). Given this, even if contextual activation of identities were occurring, it seems unlikely that this would explain our behavioral outcome.

An alternative mechanism is that introduced by Ichino and Nathan (2013), that racial (ethnic in their case) voting is moderated by expectations about racial (ethnic) targeting of goods and services. This argument makes the plausible assumption that individual voters consciously link race or ethnicity to distributive outcomes – establishing what Ferree (2006) calls a "racial heuristic". This requires political experience and knowledge, and an institutional setting that allows politicians to easily target funds to racial or ethnic groups. In particular, as Ichino and Nathan (2013) recognize, it is crucial that the political setting is competitive. For voters to change their behavior as a result of updating their beliefs about targeted spending, they must believe that there is a reasonable chance for either racial/ethnic party (or candidate) to win. As such, in non-competitive settings the rational link from context to choice is potentially broken. In many post-colonial democracies, particularly in the Southern African region, dominant parties win elections with minimal opposition, and power is highly centralized. South Africa is no exception. In these cases, it is unclear whether rational self-interest could explain the context effects we find.

A third explanation is that in-group members living in more diverse areas shift their ideology – in particular, their redistributive preferences – by coming into contact with members of an economic out-group. Given the history of South Africa, whites are economic in-group members while black South Africans are economic out-group members. Thus contact between white and black South Africans may move economic redistributive preferences of white voters toward the left – in favor of greater redistribution toward the economically needy, who are invariably black South Africans. Social psychologists have shown that emphasizing a broader identity that encompasses both the previous ingroup and outgroup can reduce ingroup bias and encourage redistribution to a disadvantaged outgroup (Gaertner, Sedikides, and Graetz 1999; Smith and Tyler 1996). In the South African context Gibson and Gouws (2005) provide evidence from just after the end of Apartheid that priming South African national identity reduces intergroup antipathy among white South Africans. Thus it is possible that white South Africans who live in more racially mixed settings come to identify more with the broader community, which in turn leads them to support redistributive policies. This third theory seems to us most plausible, given both the current state of the literature, our findings, and the contextual features of the South African case.

In speculating about the mechanisms that link racial context to racial voting in South Africa, caution should prevail. Yet we do believe that the first two mechanisms – the activation of racial expressive voting, and reasoning about distributive politics – are less plausible in this context. It seems more plausible that racial context can create shared policy preferences by shaping ideology. In this process, the preferences of whites who are more isolated become self-reinforcing, while whites exposed to other race groups update their beliefs about the plight of economically disadvantaged, and historically disenfranchised, outgroup members. Given that racial cleavages in South Africa map so closely to socio-economic cleavages, this theory certainly appeals.

4.8 Conclusion

Prior research has considered the behavioral consequences of segregation, isolation, and mixing in both developed and developing settings. Our study directly extends this work by providing a rigorous empirical examination of race context effects on political behavior in a country synonymous with both segregation and racial voting. Across three different empirical approaches we consider the effects of white isolation on white racial voting, at both the aggregate level and the individual level. These designs introduce new data sources, including high resolution Apartheid-era census data, high resolution election data for the general election of 1999, and survey data merged with geographic panel data. In each design, and at both levels of aggregation, we find consistent evidence that white isolation or segregation drives white racial voting. We do not find evidence that the results are driven by mobilization, as racial isolation fails to predict self-reported turnout intentions. Finally, we provide extensive evidence from both housing sales data and individual level survey responses to suggest that the findings are not driven by selection.

Our findings map squarely to the extant literature on South African voting behavior. Racial context effects are consistent with Ferree's (2006) synthesis of current explanations for racial voting in South Africa. While we are not in a position to adjudicate the various mechanisms that may account for our findings, we provide an elaboration of three explanations. It appears to us that the most plausible mechanism in the South African case is that inter-group racial contact encourages white voters to update their opinions about policy choices, moving them ideologically toward redistribution. As such, whites who live in more diverse settings may vote across racial lines because their policy preferences move toward the more redistributive African National Congress. Of course, other mechanisms may also play a role, and may be more dominant in settings with different histories. Future research should set out to rigorously explore and adjudicate the theoretical links we outline here.

Apartheid ranks among the most profound acts of state-led social engineering in human history, and our study provides new insights into its long term political consequences. There is a great need for empirically rigorous work on the legacies of exclusionary economic and geographic structures. Our study begins to unpack some of these consequences, and provides insights and avenues for future research. We break new ground empirically, and our newly constructed data sets provide original opportunities for social scientists to explore. Yet we would also argue the findings have bearing outside of post-Apartheid South Africa. Scholars regularly point out essential similarities between Apartheid South Africa and other countries. Racial isolation and segregation remain ubiquitous worldwide; our study shows the crucial importance of social diversity in encouraging diverse politics.

A | Appendix to Study 1: Exposure to inequality affects support for redistribution

A.1 Balance

A.1.1 Covariate Balance Between Poverty and Affluence Conditions

Table A.1: Balance in petitioner-estimated covariates across poverty (treatment) and affluence (control) conditions.

	Poverty	Affluent	Stand. diff.	P value
White	0.82	0.84	-6.02	0.12
Black	0.04	0.03	5.40	0.14
Asian	0.10	0.09	3.94	0.30
Hispanic	0.03	0.02	0.65	0.87
Old	0.12	0.11	2.28	0.56
Middle	0.44	0.44	-0.19	0.96
Young	0.44	0.45	-1.43	0.72
Female	0.57	0.57	-0.56	0.89

Table A.2: Regression of covariates on each of the eight treatment conditions. By chance, we would expect several to be statistically significant. The results suggest that there is balance in the covariates across conditions.

	Dependent variable (subject covariates):							
	Black	White	Asian	Hispanic	Young	Middle	Old	Female
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Affluent White Mill. Tax	−0.002	0.055*	−0.062**	0.001	−0.043	0.048	−0.005	0.073*
	(0.016)	(0.032)	(0.025)	(0.013)	(0.042)	(0.042)	(0.027)	(0.042)
Poor Black Mill. Tax	0.013	−0.022	−0.004	0.017	0.003	−0.011	0.005	0.037
	(0.013)	(0.026)	(0.020)	(0.011)	(0.034)	(0.034)	(0.021)	(0.034)
Poor White Mill. Tax	0.028*	−0.021	−0.018	−0.005	−0.018	0.004	0.014	0.034
	(0.015)	(0.029)	(0.023)	(0.012)	(0.039)	(0.039)	(0.025)	(0.039)
Affluent Black Placebo	−0.012	0.039	−0.024	−0.006	0.034	−0.042	0.008	0.025
	(0.013)	(0.025)	(0.019)	(0.010)	(0.033)	(0.033)	(0.021)	(0.033)
Affluent White Placebo	0.006	−0.017	−0.016	0.014	−0.016	0.034	−0.019	0.020
	(0.016)	(0.032)	(0.024)	(0.013)	(0.042)	(0.041)	(0.026)	(0.041)
Poor Black Placebo	−0.011	0.018	0.002	−0.006	0.013	−0.018	0.005	0.014
	(0.013)	(0.026)	(0.020)	(0.011)	(0.034)	(0.034)	(0.022)	(0.034)
Poor White Placebo	0.007	0.008	−0.024	−0.006	−0.037	0.036	0.001	−0.016
	(0.016)	(0.031)	(0.024)	(0.013)	(0.040)	(0.040)	(0.026)	(0.040)
Observations	2,591	2,591	2,591	2,591	2,591	2,591	2,591	2,591
R^2	0.004	0.004	0.004	0.003	0.002	0.003	0.001	0.002
Adjusted R^2	0.001	0.002	0.001	0.0004	−0.001	0.0003	−0.002	−0.001
Residual SE (df = 2583)	0.194	0.377	0.290	0.156	0.497	0.497	0.316	0.495

Note: *p<0.1; **p<0.05; ***p<0.01

A.2 Effect of Placebo Petition

Figure A.1: Plot of the difference in means for all subject in plastic bag (placebo) petition responses, at the individual level. Thin bars are 95% confidence intervals for two-sided t-tests on the difference in means; thicker bars are 80% confidence intervals. The point estimates are close to zero, and none are statistically significant at conventional levels.

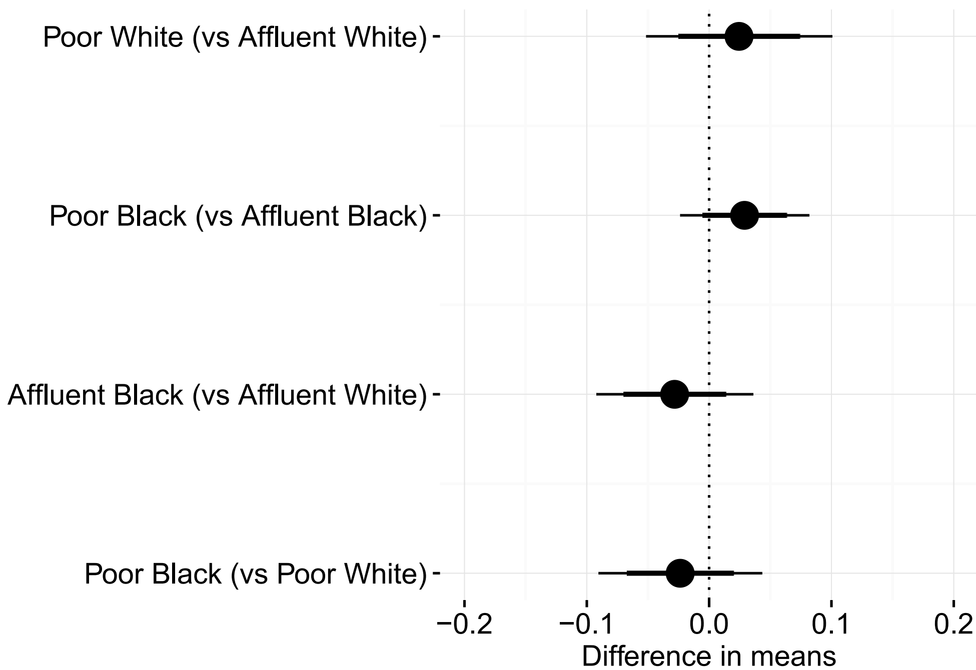

Figure A.2: Plot of the difference in means for White subjects, separated by men and women, in plastic bag (placebo) petition responses, at the individual level. Thin bars are 95% confidence intervals for two-sided t-tests on the difference in means; thicker bars are 80% confidence intervals. Most of the point estimates are close to zero, and none are statistically significant at conventional levels. Importantly, the top result (poor White versus affluent White), is null. The largest (in magnitude) point estimate corresponds to the affluent Black versus affluent White comparison; as described in the main text, this is interpreted as a racial "scare-off" effect, which is particularly strong for White subjects.

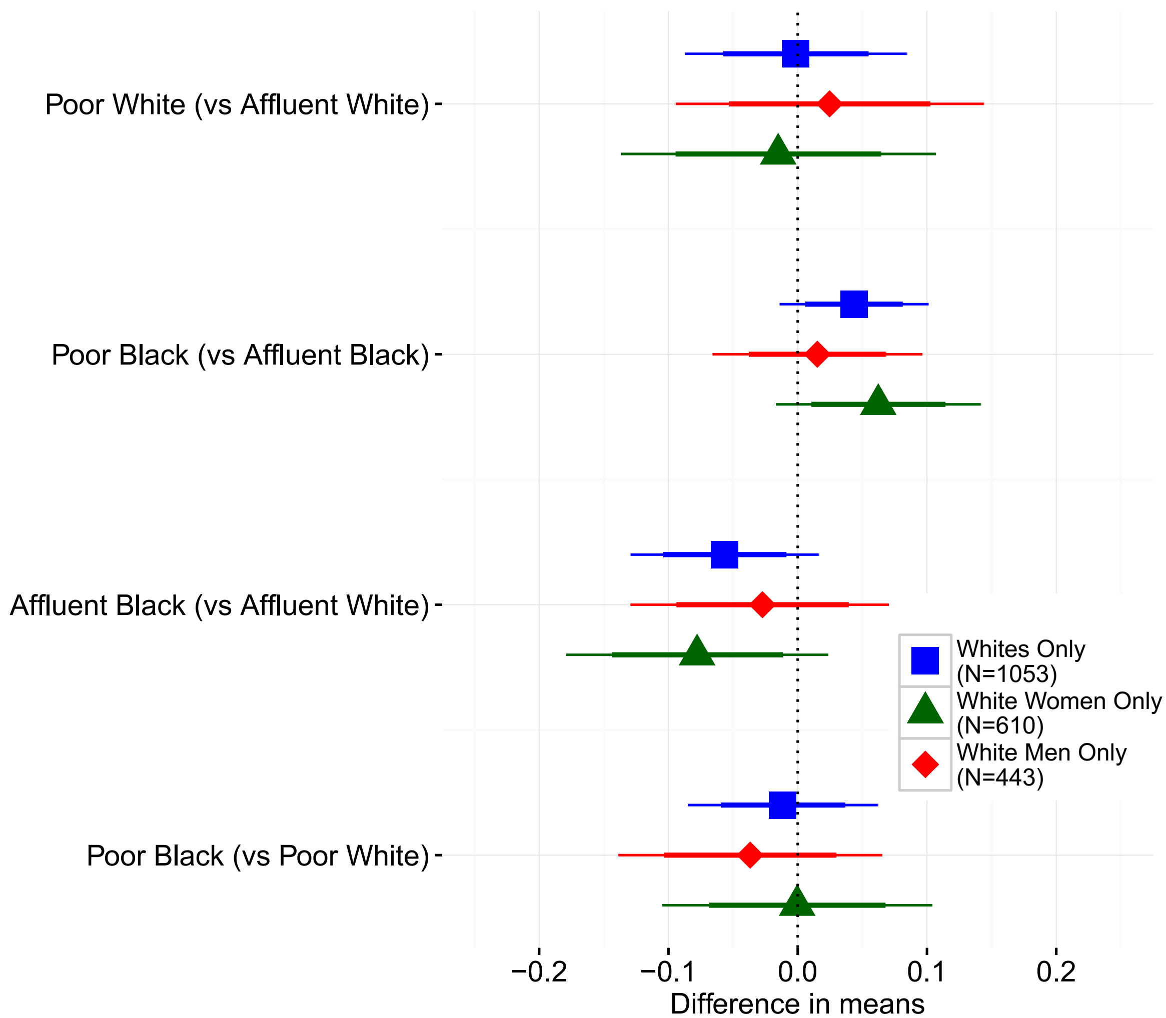

A.3 Marginal Effects Tables

Table A.3: Probit marginal effects (difference-in-differences): These specifications represent the "difference-in-differences" between the plastic bag (placebo) petition and the millionaire's tax (redistribution) petition. In columns (1) and (3), the estimated coefficients on the interaction terms correspond to the main effect of the poverty treatment. Columns (2) and (4) also include the interaction between poverty and race. Columns (3) and (4) also include controls for the petitioner's "best guess" at each subjects' race, age, and gender, as well as day fixed effects. Standard errors are adjusted for 74 day-time clusters (the level of treatment assignment). The estimated coefficients on the interaction term *Millionaire's tax* x *poor actor* in column (4) suggests that subjects are 4.2 percentage points less likely to support redistribution in the presence of an poor White person, compared to in the presence of an affluent White person ($p<0.05$). Using a linear probability model, typically the preferred approach for estimating casual effects, produces similar results. Since the probability of signing any petition is on average below 0.25, nonlinear models are shown for individual level estimates.

	Signed petition			
	(1)	(2)	(3)	(4)
Millionaire's tax	−0.085***	−0.088**	−0.088***	−0.105***
	(0.024)	(0.035)	(0.017)	(0.027)
Poor actor	0.021	0.021	0.025*	0.023*
	(0.021)	(0.021)	(0.013)	(0.013)
Black actor		−0.020		−0.044*
		(0.026)		(0.024)
Millionaire's tax x Black actor		0.005		0.023
		(0.033)		(0.026)
Millionaire's tax x Poor actor	−0.046*	−0.046*	−0.043**	−0.042**
	(0.026)	(0.026)	(0.020)	(0.020)
Day fixed effects?	No	No	Yes	Yes
Controls?	No	No	Yes	Yes
Observations	2,591	2,591	2,591	2,591
Log Likelihood	−962.417	−961.558	−936.481	−934.767

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors clustered by day-time

Table A.4: Probit marginal effects by petition: These specifications represent separate estimations for the millionaire's tax (redistribution) petition and the plastic bag (placebo) petition, respectively. Columns (2) and (4) include controls for the petitioner's "best guess" at each subjects' race, age, and gender, as well as day fixed effects. Standard errors are adjusted for 73 day-time clusters (the level of treatment assignment). The coefficient estimates on *Poor actor* in columns (1) and (2) provide evidence of a negative effect of White poverty on subjects' willingness to support redistribution, relative to White affluence. The estimates in columns (3) and (4) suggest that the apparent socioeconomic status of the confederate was not relevant to subjects' decisions about whether to support the plastic bag petition. There is some evidence, in column (4), of a Black actor "scare-off" effect; subjects were less likely to sign the placebo petition when they encountered a Black confederate. Using a linear probability model, typically the preferred approach for estimating casual effects, produces similar results. Since the probability of signing any petition is on average below 0.25, nonlinear models are shown for individual level estimates.

	Signed petition			
	Millionaire's Tax	Millionaire's Tax	Plastic Bags	Plastic Bags
	(1)	(2)	(3)	(4)
Poor actor	-0.061^{***}	-0.051^{***}	0.023	0.017
	(0.021)	(0.019)	(0.060)	(0.026)
Black actor	-0.042^{**}	-0.045^{*}	-0.029	-0.067^{*}
	(0.021)	(0.023)	(0.053)	(0.037)
Poor actor x	0.071^{**}	0.056^{**}	0.007	0.023
Black actor	(0.035)	(0.028)	(0.067)	(0.032)
Day fixed effects?	No	Yes	No	Yes
Controls?	No	Yes	No	Yes
Observations	1,335	1,335	1,256	1,256
Log Likelihood	-356.952	-345.454	-602.430	-572.613
Akaike Inf. Crit.	721.904	738.907	1,212.861	1,193.226

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
Standard errors clustered by day-time

A.4 Difference-in-differences Cluster Level Regression

Table A.5: Difference-in-difference cluster-level treatment effects: These specifications represent the "difference-in-differences" between the plastic bag (placebo) petition and the millionaire's tax (redistribution) petition at the cluster level. Based on column 1, subjects are 8 percentage points *less* likely to support the redistributive policy in the presence of a poor person, net of baseline response rates ($p<0.05$). Column 2 indicates that this effect is driven exclusively by the poor White actor.

	Signed petition	
	(1)	(2)
Millionaire's tax	-0.067^{**}	-0.079^{**}
	(0.026)	(0.035)
Poor actor	0.036	0.010
	(0.027)	(0.035)
Black actor		-0.055
		(0.034)
Poor actor x Black actor		0.042
		(0.038)
Millionaire's tax x Poor actor	-0.080^{**}	-0.076^{**}
	(0.038)	(0.038)
Millionaire's tax x Black actor		0.017
		(0.038)
Constant	0.180^{***}	0.214^{***}
	(0.019)	(0.028)
Clusters (observations)	74	74
Individuals	2591	2591
Residual SE	0.081 (df = 70)	0.081 (df = 67)
Note:	*p<0.1; **p<0.05; ***p<0.01	

A.5 Photographs of Experiment in Progress

Figure A.3: Petitioner with confederate in Brookline, MA

Figure A.4: Petitioner with confederate in Boston, MA

A.6 Experiment Materials

Script for Petitioners

Would you sign this petition in support of the Millionaire's Tax in Massachusetts?

If asked for more details:

I'm gathering signatures to send a message to the Massachusetts legislature urging them to support a tax increase on earnings over one million dollars.

Would you like to see the text of the petition? [Show petition]

OR

Would you sign this petition in support of reducing wasteful plastic bags in Massachusetts?

If asked for more details:

I'm gathering signatures to send a message to the Massachusetts legislature urging them to pass a law that would encourage people to reduce the use of disposable plastic bags.

Would you like to see the text of the petition? [Show petition]

Figure A.5: Petitioners used this tracking sheet to unobtrusively record their best guesses of subject characteristics, including those who do and do not sign.

Tracking Sheet Date:___________________________

Sheet #:_______ of ________

Signed petition?		Refusal type				Characteristics			Time
Mill.'s Tax	Plastic Bags	Polite	Reasoned	Evasive	Hostile	Gender M/F	Age Y/M/O	Race W/B/A/H	HH:MM

Mark with "U" if unsolicited Y: 18-35; M: 35-65; O: 65+

A.6.1 Randomization

The randomization occurred at the start of each day (block) such that all subjects had an equal chance of being assigned to one of the eight treatment conditions. Each day, three features of the starting condition were randomized independently: (1) the petition (millionaire's tax or plastic bags); (2) the race of the actor (phenotypically Black or White); and (3) the actor's role (poor or affluent). The petition was switched halfway through the day, and the actors switched on and off every thirty minutes to minimize the gaps between conditions.

A.7 Inequality vs. Poverty

Economic inequality is an abstract concept, difficult to concretely portray without the help of numbers, graphs, or words. To buttress the claim that placing poor confederates in an affluent neighborhood signifies *inequality* – distinct from, or rather than exclusively, poverty – I conducted an online survey experiment on Amazon's Mechanical Turk (n = 118, after dropping subjects who failed the attention check). In it, subjects responded to photographs of scenes like those staged in the field, as well as scenes of poor people without environmental features of affluence, and affluent people without signs of poverty. One example photo is shown below.

After being randomly assigned to view a photo of inequality, poverty, or affluence, subjects were asked to describe those images in a word or words. They were then shown the same photo a second time, along with a set of five words ("wealth", "poverty", "middle class", "inequality" and "tall"), and were asked "which of the following words best describes this image? Rank the words from 1 (most describes the image) to 5 (least describes the image)."

The results provide compelling evidence that the experimental treatment is exposure to inequality rather than to poverty. Without being primed by words related to inequality, subjects who viewed figure A.6 and similar images gave responses that included "inequality", "unequal", "disparity", "income gap", "the haves versus have-nots", and "social injustice" when asked to describe the images in an open-ended text field. Though respondents used a plethora of different descriptive words or phrases, the modal answer reflected inequality or disparity. Meanwhile the modal response to the poverty only photos was "sad". In the next question, subjects ranked "inequality" as the word that most aptly described photos of poor people in wealthy environments. As shown in table A.6, the median rank of "inequality" was 1 (most describes the image), and the mean 1.8 (out of 5). The last two rows of table A.6 show that 64% of respondents ranked "inequality" above "poverty" when describing an image of a poor person in an affluent setting, while only 4% of respondents shown an image of just poor people (e.g., figure A.7) ranked “inequality” above “poverty”.

Table A.6: Summary statistics for photograph ranking exercise on mTurk. “Inequality condition” refers to subjects randomly assigned to view a photo of inequality (e.g., figure A.6); “poverty condition” refers to subjects randomly assigned to view a photo of poverty (e.g., figure A.7). Subjects were given a set of five words (“wealth”, “poverty”, “middle class”, “inequality” and “tall”) and asked rank the words from 1 (most describes the image) to 5 (least describes the image).

	mean	sd	median	min	max	n
Inequality rank - inequality condition	1.84	1.26	1.00	1.00	5.00	45
Poverty rank - inequality condition	2.36	1.11	2.00	1.00	5.00	45
Wealth rank - inequality condition	2.67	0.88	3.00	1.00	4.00	45
Ranked inequality above poverty - inequality cond.	0.64	0.48	1.00	0.00	1.00	45
Ranked inequality above poverty - poverty cond.	0.04	0.20	0.00	0.00	1.00	48

As additional evidence that the field experiment treatment approximates exposure to inequality, note that the results of a Google Image search for “images of inequality”, restricting the search to photos, yields photos of poor or homeless individuals in close proximity to more affluent individuals. This suggests that the experimental setup reflects popular conceptions of income inequality as a real-world phenomenon.

Figure A.6: Example of photo portraying inequality.

Figure A.7: Poverty, without affluence

Figure A.8: Affluence, without poverty

B | Appendix to Study 2: 'Eyes' on the Street

B.1 Census Block Crime and Foot Traffic

Table B.1: Relationship between major crime incidents per capita and mean foot traffic at the Census block level. Blocks with zero residential population are dropped from the analysis. Columns 2 and 3 include controls for the racial and ethnic makeup of Census blocks; column 3 also contains Census tract fixed effects.

	Crime incidents per capita		
	(1)	(2)	(3)
Foot traffic	0.0001^{***}	0.0001^{***}	-0.00003
	(0.00003)	(0.00003)	(0.0001)
Constant	-0.107^{***}	-0.079^{**}	0.058
	(0.029)	(0.034)	(0.089)
Controls?	No	Yes	Yes
Fixed effects?	No	No	Yes
Observations	512	512	512
Adjusted R^2	0.037	0.039	0.196
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$		

C | Appendix to Study 3: Can violent protest change local policy support?

C.1 Tests for Similar Trends in Distance Relationship Prior to the Riot

As an additional robustness check, we look to see if similar relationships between distance and support for public schools can be found prior to the riot. We use pairs of parallel ballot initiatives in which both votes occur *prior* to the riots. In both the 1986 and 1990 general elections in California, there were ballot initiatives proposing bonds for public schools and higher education.

Initiative	Year	Title	Dollar Amount	Result
53	1986	Public School Construction Bonds	$800M	Pass
54	1986	Bonds for Higher Education Facilities	$400 M	Pass
143	1990	Public School Construction Bonds	$800 M	Fail
146	1990	Bonds for Higher Education	$450 M	Pass

Table C.1: Two pairs of parallel ballot initiatives were placed before voters in 1986 and 1990, all prior to the Los Angeles Riots.

We construct the same difference-in-differences estimator as before, using different ballot initiatives:

$$\texttt{EdDiffPlacebo}_i = (\texttt{PubSchool}_{1990i} - \texttt{PubSchool}_{1986i}) - (\texttt{HigherED}_{1990i} - \texttt{HigherED}_{1986i})$$

We expect that since there is no event prior and proximate to the 1990 general election to affect the salience of distance from that particular location, a regression of distance from Florence and Normandie on `EdDiffPlacebo` will be zero or negative, as compared to the positive effect we expected from the true test.

	(1)	(2)
Distance	0.013*	−0.047*
	(0.001)	(0.004)
Distance^2		0.014*
		(0.001)
Constant	0.067*	0.111*
	(0.002)	(0.003)
N	2,560	2,560
R^2	0.04	0.09

Table C.2: OLS regression of `EdDiffPlacebo` and distance from Florence and Normandie, measured in units of 10 kilometers. Standard errors are in parentheses. $*$ represents $p < .05$.

As shown in table C.2, the relationship between `EdDiffPlacebo` and distance from the future location of the riot is, for white voters, smaller and in the opposite direction of the what we observed in Table 2. This final test suggests that, in the absence of the riots, white voting behavior may have been markedly different as a function of distance from Florence and Normandie.

C.2 Survey Data Analysis

Table C.3 is the covariate balance on white LACSS respondents before and after the riot.

Table C.4 is the covariate balance on white LACSS respondents before and after the riot.

In Table C.5 we show regressions on support for spending on "improving our nation's education system". The independent variable of interest is being asked either before or after the verdict. Both excluding (column 1) and including demographic controls (column 2), this coefficient is near zero. As shown in columns 3 and 4, we see no evidence for changes in

	Before riot	After riot	Difference	p value
High School	0.91	0.90	-0.01	0.61
Attended College	0.66	0.72	0.06	0.20
Income $< 30k$	0.54	0.46	-0.08	0.13
Female	0.67	0.57	-0.10	0.04
Homeowner	0.38	0.50	0.12	0.01
Married	0.34	0.32	-0.02	0.67
Unemployed	0.03	0.04	0.01	0.63
Age	44.91	43.48	-1.43	0.41
Conservative	0.15	0.11	-0.04	0.25
Democrat	0.73	0.72	-0.01	0.85
Republican	0.01	0.05	0.03	0.04
Independent	0.17	0.16	-0.02	0.68
Distance from Florence/Normandie	3092	3125	33	0.91

Table C.3: Covariate balance in LACSS sample, African American respondents. All values represent proportions, except for age (years) and distance (meters). P values generated by T-test for difference of means (n = 426). An omnibus balance test provides some evidence that we cannot reject the null hypothesis that the data are balanced (p=0.132, χ^2 test statistic with 16 degrees of freedom = 22.4).

	Before riot	After riot	Difference	p value
High School Grad	0.92	0.95	0.03	0.45
Attended College	0.84	0.82	-0.02	0.76
Income $< 30k$	0.35	0.36	0.01	0.85
Female	0.46	0.55	0.09	0.25
Homeowner	0.48	0.64	0.15	0.04
Married	0.38	0.49	0.11	0.17
unemployed	0.02	0.05	0.03	0.35
Age	44.43	50.30	5.87	0.03
Conservative	0.17	0.21	0.04	0.47
Democrat	0.40	0.40	-0.00	0.95
Republican	0.25	0.26	0.02	0.82
Independent	0.24	0.28	0.04	0.56
Distance (km)	11320	8893	-2427	0.00

Table C.4: Covariate balance in LACSS sample, white respondents. All values represent proportions, except for age (years) and distance (meters). P values generated by T-test for difference of means (n = 185). In an omnibus balance test we reject the null hypothesis that the data are balanced (p=0.002, χ^2 test statistic with 16 degrees of freedom = 37.2)

ideology either.

	Spending too little on education		More conservative	
	(1)	(2)	(3)	(4)
After verdict	−0.00	0.02	0.00	−0.03
	(0.04)	(0.04)	(0.09)	(0.10)
Constant	2.89*	2.83*	1.80*	1.61*
	(0.03)	(0.08)	(0.07)	(0.21)
Controls?	No	Yes	No	Yes
N	418	380	272	248
R^2	0.00	0.08	0.00	0.03

Table C.5: OLS regression of measures of attitude change on an indicator for whether respondents were surveyed after the Rodney King verdict. African American LACSS respondents only. "After Verdict" is an indicator variable for whether respondents were interviewed before or after the announcement of the verdict in the trial of the police officers. The dependent variable in Columns 1 and 2 is support for education spending, as measured by the degree to which respondents agree that too little is spent to improve education (respondents indicated whether we are "spending too much", "spending the right amount" or "spending too little" on "improving the nation's educational system"). The dependent variable in Columns 3 and 4 is ideology, "Would you consider yourself as a... Liberal, Conservative, Moderate, or don't you consider yourself that way?" Sample size decrease in the latter columns is due to respondents who claimed to not consider themselves that way. Control variables include respondent age, home ownership, marital status, gender, education, income, and distance from Florence and Normandie. Standard errors are in parentheses. $*$ represents $p < .05$.

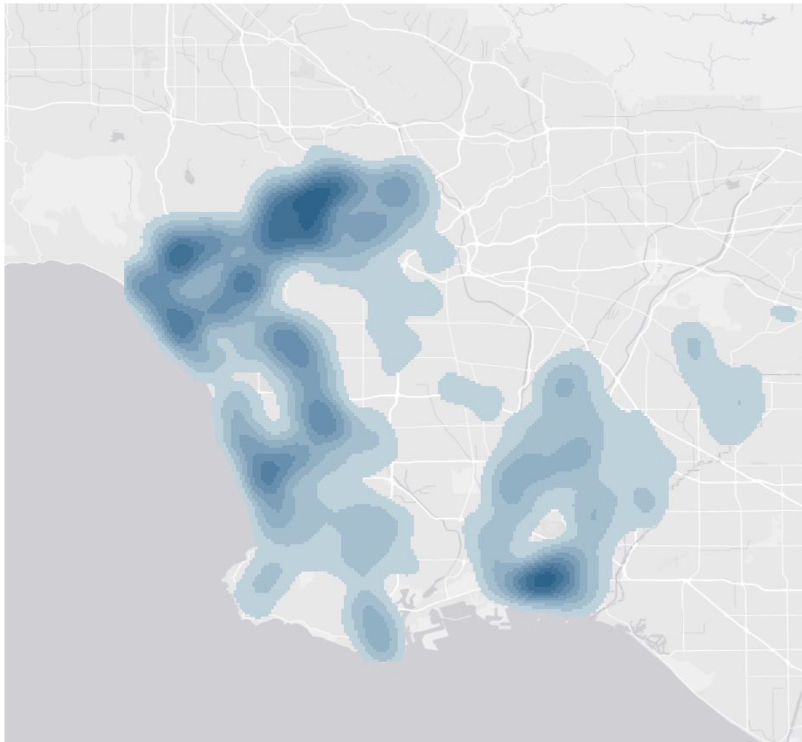

Figure C.1: Kernel density map of whites who registered to vote in the 5 weekdays prior to the riot

C.3 Distribution of Voters Registering Before and After the Riot

Figures C.1–C.4 are kernal-density maps of the distribution of white and African American registrants before and after the riot.

C.4 Voter File Matching and Attrition

In matching the 1992 and 2005 voterfiles, we first match exactly on first name, last name, and date of birth. Then, among the remaining women only, we match on first name, middle name, and date of birth. This second merge is designed to capture women who changed their last names, and in practice matches few additional registrants. Because we only examine a

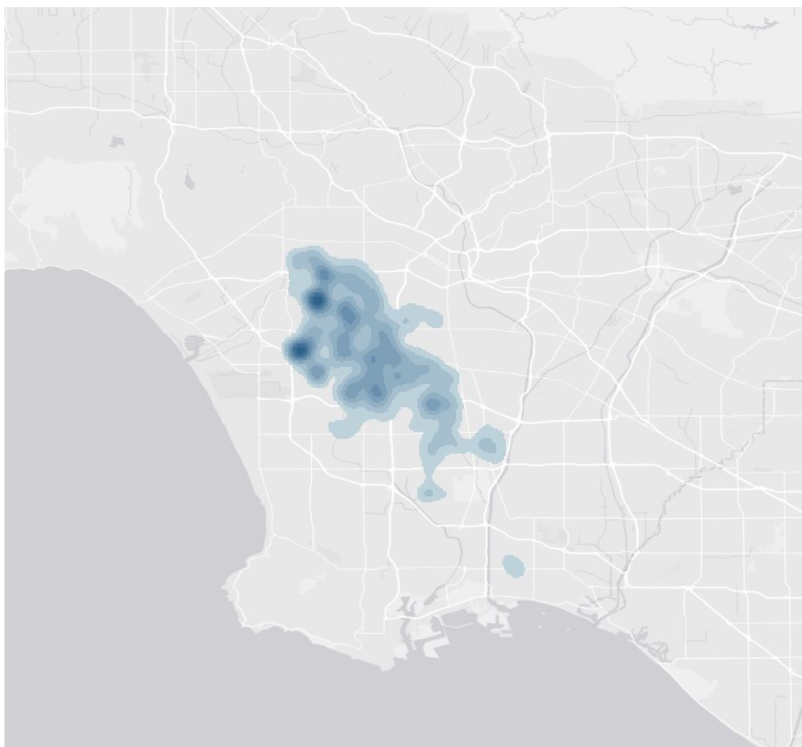

Figure C.2: Kernel density map of African Americans who registered to vote in the 5 weekdays prior to the riot

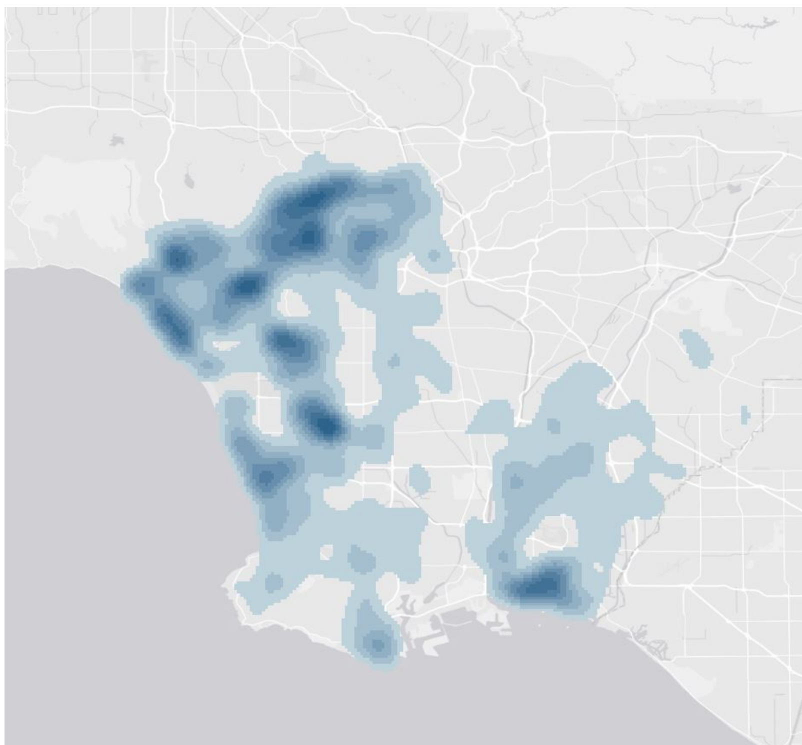

Figure C.3: Kernel density map of whites who registered to vote in the 5 weekdays following the riot

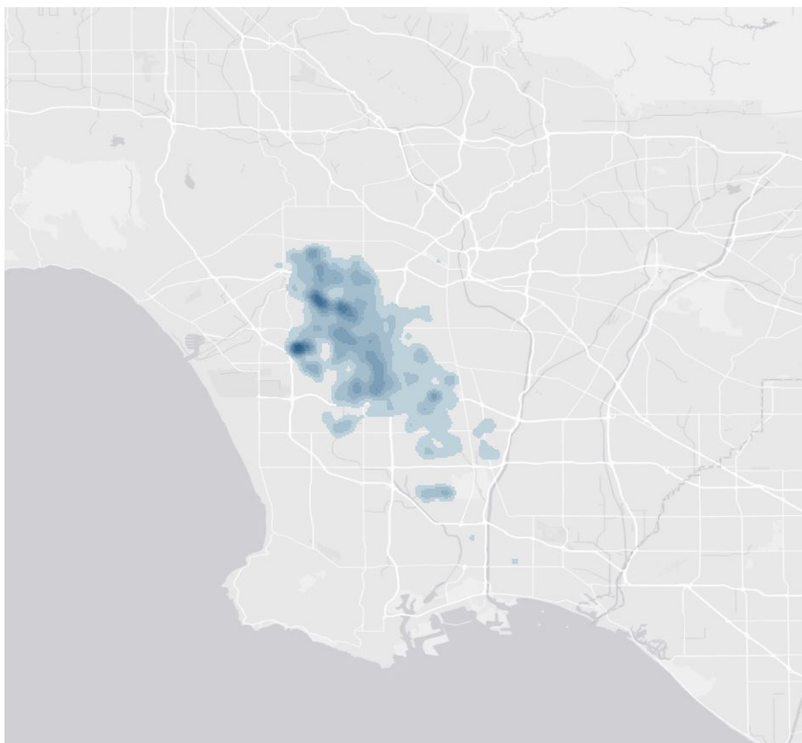

Figure C.4: Kernel density map of African Americans who registered to vote in the 5 weekdays following the riot

California file, we necessarily cannot find voters who left the state. We also lose voters who change their name due to marriage. Other voters die, or move and then do not re-register, and are subsequently removed from the voter file due to inactivity.

Here we detail calculations used to estimate the percentage of the 1992 Voter File we expect to be able to identify in 2005. In all cases, we aim to produce conservative estimates which will underestimate the number of people we expect to have attritted. Overall, we expect more than 9,600 individuals from our 1992 sample of 30,166 to have attritted (see below). In step 1, we match exactly on first name, last name, and date of birth. In this stage, we locate 19,165 individuals. In the second stage, we merge the remainder on first name, middle name, and date of birth among women only. Here, we locate an additional 1,920 individuals, bringing our merged total to 15,244. While we estimate that about 3,314 women in our original sample will have changed their names due to marriage, we are only able to locate 1,920.

We tried several different alternate merge schemes, including matching last names from 1992 to "previous last name" from the 2005 file along with first name and date of birth, but this produced no matches. Additionally, we tried a fuzzy merge where we merged within a set Levenshtein distance from a merge identifier consisting of first name, last name, date of birth, and imputed race, but this produced an excess of false-positive matches.

In order to assess the performance of our matching exercise, we estimate the total amount of expected attrition from the voter file due to (1) name changes at marriage, (2) death, and (3) migration.

Marriage To calculate marriage, we use statistics on marriage rates by age among women as of year 2000. Since marriage rates have been decreasing, we believe this should produce an underestimate of the true number of married women who have changed their names.

Using date of birth from the voter file, we calculate the age of everyone in the sample in year 1992. The Bureau of Labor Statistics provides marriage probabilities binned by age: 25% of women under the age of 25 are married; 53% between 26 and 35; 81% between 36 and 45; and 86% of women aged 46 and over. Using this formula, we calculate the number of women married in 1992, and the number of the unmarried women who would become married by 2005, and subtract those numbers. Finally, I multiply this difference by 0.85, corresponding to the 15% of women who do not change their name at all after marriage. This totals to 2,457 women whom we estimate changed their last names. While in theory we may be able to find these women in the second stage of our merge, in practice our voter file often excludes middle name, resulting in failed merges for this stage.

Death According to the CDC, the yearly mortality rate for African Americans is 733 per 100,000; for white and Hispanic people, it is 350 per 100,000 (CDC 2016). Using our imputed race values, and accounting for yearly compound reductions in population, we estimate a sample attrition of 824 people.

Migration Out-migration from California as averaged 500,000 per year since 1990 (Perry et al. 2016). Out-migration has disproportionately happened from urban areas; we assume that this out-migration has happened equally across the state, which should produce an underestimate of our sample decay.

California contained approximately 31 million people in 1990. We multiply $31,0166 * (500,000/31,000,000) * 13 = 6,325$ individuals from our sample who we expect to have left California, and thus to no longer appear on the voter rolls.

Total Summing these components, we estimate attrition of 9,606 individuals between 1992 and 2005, of whom 2,457 are potentially identifiable using voter file variables which involve previous registration names. Aside from this, there are additional reasons why people may be invalidated from voting, including felony convictions, incarceration, and mental illness.

	White		African American	
	(1)	(2)	(3)	(4)
(Intercept)	0.94*	0.92*	0.91*	0.88*
	(0.01)	(0.01)	(0.02)	(0.04)
Registered Post-Riot	0.00	0.00	0.01	0.02
	(0.01)	(0.02)	(0.01)	(0.02)
Controls	No	Yes	No	Yes
N	6,887	3,294	3,471	1,526
R^2	0.00	0.00	0.00	0.01

Table C.6: Turnout in 2004 general election for those registering immediately before and immediately after the riot. OLS regression on voter turnout. Registered post-riot is a dummy variable indicating those who registered before and after the riot. Controls include age, gender, party ID. Standard errors are in parentheses. $*$ represents $p < .05$.

	White		African American	
	(1)	(2)	(3)	(4)
(Intercept)	0.24*	0.20*	0.23*	0.28*
	(0.01)	(0.02)	(0.03)	(0.05)
Registered Post-Riot	-0.01	-0.02	0.00	-0.04
	(0.01)	(0.02)	(0.02)	(0.03)
Controls	No	Yes	No	Yes
N	6,887	3,294	3,471	1,526
R^2	0.00	0.00	0.00	0.01

Table C.7: Turnout in 2004 primary election for those registering immediately before and immediately after the riot. OLS regression on voter turnout. $*$ represents $p < .05$.

Finally, while more than 7,000 voters became *ineligible* to vote due to death or outmigration, it is an open question as to how many changed their address and failed to re-register, or were purged from voter rolls due to inactivity. Overall, we expect that our merging procedure has approximately captured a sufficient sample of the registrants from our 1992 population who were still registered in 2005.

C.5 Long-term Participation of Riot Registrants

In Tables C.7 and C.6 we regress registration before and after the riot on participation in the 2004 General and Primary elections, respsctively.

D | Appendix to Study 4: Segregation drives racial voting

D.1 Variable Definitions and Summary Statistics

Table D.1: Ward-level variables and definitions

Variable	Definition
WhiteIsolation	The proportion of white people who reside in the neighborhood of an average white person. Higher values indicated higher levels of white-centric segregation
BlackFrac	The count of black individuals divided by the total population
ColFrac	The count of coloured individuals divided by the total population
WhiteFrac	The count of white individuals divided by the total population
Income	The average income, from a 17-point discrete scale, logged
Education	The average level of education, from a 7-point discrete scale
$Employ_B$	The number of employed individuals divided by the total working age population
$Employ_N$	The number of employed individuals divided by the total work-seeking population
ANCvs2014	The total number of votes cast for the ANC in the 2014 National ballot, divided by the total number of votes cast
ANCvs1999	The total number of votes cast for the ANC in the 1999 National ballot, divided by the total number of votes cast

Table D.2: Ward-level summary statistics

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
WhiteIsolation1991	0.322	0.372	0	1	2920
WhiteIsolation2011	0.159	0.248	0	0.943	4277
BlackFrac1991	0.629	0.404	0	1	2919
BlackFrac2011	0.814	0.305	0.011	1	4277
ColFrac1991	0.144	0.296	0	1	2919
ColFrac2011	0.09	0.223	0	0.982	4277
WhiteFrac1991	0.2	0.299	0	1	2919
WhiteFrac2011	0.077	0.173	0	0.940	4277
Income1991	7.944	1.191	-0.2	11.026	2905
Income2011	7.16	0.823	-0.254	10.097	4276
Education1991	0.432	0.105	0	0.833	2919
Education2011	0.497	0.083	0	0.792	4277
$Employ_B 1991$	0.317	0.181	0	1	2919
$Employ_B 2011$	0.349	0.176	0	0.933	4277
$Employ_N 1991$	0.797	0.193	0	1	2904
$Employ_N 2011$	0.642	0.166	0.197	0.989	4275
ANCvs2014	0.642	0.256	0.005	0.983	4268
ANCvs1999	0.666	0.302	0.002	0.991	4138

Table D.3: Individual-level variables and definitions

Variable	Definition
Retro ANC Vote	Binary indicator of whether vote in last election was for the ANC
Future ANC Vote	Binary indicator of intention to vote for the ANC in future election
Turnout	Binary indicator of whether voted in last election
Coloured	Binary indicator for membership of Coloured race group
Indian	Binary indicator for membership of Indian/Asian race group
White	Binary indicator for membership of White race group
High School	Binary indicator for completion of high school (grade 12)
Prim School	Binary indicator for completion of primary school (grade 7)
Married	Binary indicator for whether married
Sex	Binary indicator for gender
Age	Age in completed years
Age Sq	Age squared

Table D.4: Individual-level summary statistics

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
Retro ANC Vote	0.554	0.497	0	1	35267
Future ANC Vote	0.522	0.5	0	1	39168
Turnout	0.748	0.434	0	1	35267
Coloured	0.161	0.368	0	1	39168
Indian	0.103	0.304	0	1	39168
White	0.124	0.33	0	1	39168
Wealth	0.622	0.372	0	1	35916
High School	0.393	0.488	0	1	39168
Prim School	0.813	0.39	0	1	39168
Married	0.554	0.497	0	1	38845
Sex	0.593	0.491	0	1	39168
Age	39.869	16.15	16	99	39168
Age Sq	1850.307	1461.186	256	9801	39168

D.2 First Stage Instrumental Variables Results

Table D.5: First stage relationship

	Dependent variable:			
	WhiteIsolation2011			
	(1)	(2)	(3)	(4)
Instruments				
log_mean_tri	0.0326***		0.0290***	
	(0.0101)		(0.0105)	
nn_lra		0.0378***		0.0304***
		(0.00773)		(0.00836)
Covariates				
white_iso1991	0.312***	0.302***	0.289***	0.280***
	(0.0277)	(0.0287)	(0.0275)	(0.0286)
white_frac1991	0.391***	0.399***	0.429***	0.429***
	(0.0470)	(0.0465)	(0.0414)	(0.0411)
black_frac1991	0.0524*	0.0478*	0.0543*	0.0597**
	(0.0284)	(0.0281)	(0.0287)	(0.0284)
colored_frac1991	0.0265	0.0341	0.0334	0.0459
	(0.0375)	(0.0365)	(0.0362)	(0.0357)
Observations	2,873	2,905	2,857	2,889
R^2	0.714	0.716	0.731	0.730
Municipal FE	✔	✔	✔	✔
Extra covariates			✔	✔

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

D.3 Evidence Against Sorting and Validity of Instrumental Variables Analysis

Table D.6: Correlations between covariates and instruments over time. Statistical tests for the equivalence of the correlation coefficients suggest that none of the coefficients are statistically distinguishable between time periods.

	Instrument:	
	LogMeanTRI	CNLRA
Income1991	-0.0752	-0.2379
Income2011	-0.0321	-0.1935
Educ1991	-0.1024	-0.3911
Educ2011	-0.1344	-0.4253
Employ_N1991	-0.1496	-0.1676
Employ_N2011	-0.0728	-0.1094
Employ_B1991	-0.1129	-0.2114
Employ_B2011	-0.1289	-0.2296

Table D.7: Ignorability of Instruments. We estimate the “first stage" relationship of each instrument on four economic covariates in the baseline 1991 data.

	Dependent Variable:							
	log_income1991	log_income1991	educ1991	educ1991	employ_n1991	employ_n1991	employ_b1991	employ_b1991
log_mean_tri	-0.0283		-0.0107***		-0.00817		0.000378	
	(0.0368)		(0.00394)		(0.00728)		(0.00771)	
nn_lra		-0.0553**		-0.0256***		0.00418		0.000305
		(0.0278)		(0.00437)		(0.00602)		(0.00576)
white_iso1991	0.566***	0.603***	0.00369	0.0121**	0.0488***	0.0508***	0.0555***	0.0602***
	(0.0706)	(0.0701)	(0.00630)	(0.00573)	(0.0122)	(0.0124)	(0.0176)	(0.0177)
black_frac1991	-1.010***	-0.989***	-0.0951***	-0.0912***	-0.154***	-0.156***	-0.0244	-0.0222
	(0.110)	(0.105)	(0.0128)	(0.0130)	(0.0327)	(0.0323)	(0.0369)	(0.0363)
white_frac1991	0.658***	0.643***	0.115***	0.113***	0.0543*	0.0492	0.0978***	0.0950**
	(0.131)	(0.127)	(0.0141)	(0.0141)	(0.0301)	(0.0301)	(0.0374)	(0.0373)
colored_frac1991	-0.632***	-0.637***	-0.0632***	-0.0660***	-0.109***	-0.112***	-0.0104	-0.0124
	(0.126)	(0.122)	(0.0155)	(0.0165)	(0.0333)	(0.0340)	(0.0336)	(0.0337)
Observations	2,859	2,891	2,873	2,905	2,858	2,890	2,873	2,905
R^2	0.725	0.735	0.715	0.738	0.548	0.550	0.549	0.550
Municipal FE	✔	✔	✔	✔	✔	✔	✔	✔

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table D.8: Regression of deeds transfers per person on 1991 EA-level covariates. Column (1) suggests that, at the neighborhood level, there is no relationship between sorting (proxied by property deed transfers) and the White proportion of the population. Columns (2) and (3) show that there is also no relationship between sorting and the proportion of residents who speak Afrikaans, or other characteristics of the neighborhood, respectively. "White flight" does not appear to be a significant determinant of real estate sales in the 1990s.

	Dependent variable:		
	Deeds transfers per person (1993-1999)		
	(1)	(2)	(3)
whitefrac	−0.007	−0.004	0.003
	(0.007)	(0.009)	(0.010)
coloredfrac	−0.020	−0.016	−0.008
	(0.014)	(0.016)	(0.016)
asianfrac	−0.021	−0.021	−0.011
	(0.018)	(0.018)	(0.019)
afrikaans		−0.006	−0.015
		(0.010)	(0.011)
household_income			−0.00000
			(0.00000)
population_density			−0.00000***
			(0.00000)
property_owners			−0.046***
			(0.016)
Constant	0.022***	0.022***	0.032***
	(0.005)	(0.005)	(0.006)
Observations	4,509	4,509	4,509
R^2	0.001	0.001	0.004
Adjusted R^2	0.0001	−0.0001	0.003

*p<0.1; **p<0.05; ***p<0.01

Table D.9: Regression of segregation in 2011 on hilliness, deeds transfers (1993-1999), and baseline (1991) covariates. Controlling for deeds transfers, a proxy for sorting, does not affect the first stage estimates.

	Dependent variable:			
	white_iso2011			
	(1)	(2)	(3)	(4)
Deeds transfers per person		-0.003^{***}		-0.003^{***}
		(0.00001)		(0.00000)
log_mean_tri	0.051^{***}	0.049^{***}		
	(-0.0004)	(-0.0004)		
nn_lra			0.084^{***}	0.085^{***}
			(-0.001)	(-0.001)
Observations	595	595	608	608
R^2	0.646	0.648	0.658	0.661
Adjusted R^2	0.628	0.630	0.641	0.643
Municipal FE	✔	✔	✔	✔
Baseline Covs	✔	✔	✔	✔

Note: Cluster robust standard errors in parentheses *p<0.1; **p<0.05; ***p<0.01

Table D.10: Regression of deeds transfers per person on the hilliness instruments and 1991 baseline characteristics. The large standard errors suggest that hilliness does not predict sorting.

	Dependent variable:	
	Deeds transfers per person (1993-1999)	
	(1)	(2)
log_mean_tri	−0.797	
	(−2.911)	
nn_lra		−0.548
		(−1.336)
Observations	595	608
R^2	0.105	0.106
Adjusted R^2	0.054	0.057
Municipal FE	✔	✔
Baseline Covs	✔	✔
Note: Cluster robust standard errors in parentheses		*p<0.1; **p<0.05; ***p<0.01

Table D.11: Individual interaction results for measures of racial sentiment

	Dependent variable:			
	Pos. to White	Friendly to White	Pos. to Black	Friendly to Black
white	0.0864**	0.119***	-0.107***	-0.111***
	(0.0394)	(0.0373)	(0.0389)	(0.0416)
white_iso2011	0.0227	0.0126	-0.00119	0.0836**
	(0.0414)	(0.0364)	(0.0322)	(0.0324)
white_white_iso	0.0127	-0.00341	0.0118	-0.0120
	(0.0644)	(0.0635)	(0.0659)	(0.0670)
N	4,677	4,682	4,699	4,704
R^2	0.111	0.121	0.129	0.141
Indiv Covariates	✔	✔	✔	✔
Ward Covariates	✔	✔	✔	✔
Municipal FE	✔	✔	✔	✔

Cluster robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Bibliography

Abrajano, Marisa A, and R Michael Alvarez. 2005. "A natural experiment of race-based and issue voting: The 2001 city of Los Angeles elections." *Political Research Quarterly* 58 (2): 203–218.

Abrajano, Marisa A, R Michael Alvarez, and Jonathan Nagler. 2008. "The Hispanic vote in the 2004 presidential election: Insecurity and moral concerns." *The Journal of Politics* 70 (02): 368–382.

Acharya, Avidit, Matthew Blackwell, and Maya Sen. 2016. "The Political Legacy of American Slavery." *Journal of Politics* 78 (3): 621–641.

Adida, Claire L, Nathan Combes, Adeline Lo, and Alex Verink. 2016. "The Spousal Bump Do Cross-Ethnic Marriages Increase Political Support in Multiethnic Democracies?" *Comparative Political Studies*: 0010414015621080.

Agnew, John. 1996a. "Mapping politics: how context counts in electoral geography." *Political Geography* 15 (2): 129–146.

Agnew, John. 1996b. "Maps and models in political studies: a reply to comments." *Political Geography* 15 (2): 165–167.

Alesina, Alberto, and Ekaterina Zhuravskaya. 2008. Segregation and the Quality of Government in a Cross-section of Countries. Technical report National Bureau of Economic Research.

Alesina, Alberto, Reza Baqir, and William Easterly. 1999. "Public Goods and Ethnic Divisions." *The Quarterly Journal of Economics* 114 (4): 1243–1284.

Allport, Gordon Willard. 1979. *The nature of prejudice.* Basic Books.

Alston, Jon P, and K Imogene Dean. 1972. "Socioeconomic factors associated with attitudes toward welfare recipients and the causes of poverty." *The Social Service Review*: 13–23.

Alvarez, R Michael, Stephen Ansolabehere, Adam Berinsky, Gabriel Lenz, Charles Stewart III, and Thad Hall. 2009. "2008 Survey of the Performance of American Elections: Final Report." *Manuscript, Caltech/MIT Voting Technology Project. http://dspace. mit. edu/bitstream/handle/1721.1/49847/Final% 20report20090218. pdf.*

Ananat, Elizabeth Oltmans. 2011. "The wrong side(s) of the tracks: The causal effects of racial segregation on urban poverty and inequality." *American Economic Journal: Applied Economics*: 34–66.

Ananat, Elizabeth Oltmans, and Ebonya Washington. 2009. "Segregation and Black political efficacy." *Journal of Public Economics* 93 (5): 807–822.

Bajari, Patrick, and Matthew E Kahn. 2005. "Estimating housing demand with an application to explaining racial segregation in cities." *Journal of Business & Economic Statistics* 23 (1).

Bartels, Larry M. 2009. *Unequal democracy: The political economy of the new gilded age.* Princeton University Press.

Bates, Robert H. 1983. "Modernization, ethnic competition, and the rationality of politics in contemporary Africa." *State versus ethnic claims: African policy dilemmas* 152: 171.

Bayer, Patrick, Robert McMillan, and Kim S Rueben. 2004. "What drives racial segregation? New evidence using census microdata." *Journal of Urban Economics* 56 (3): 514–535.

Beber, Bernd, Philip Roessler, and Alexandra Scacco. 2014. "Intergroup Violence and Political Attitudes: Evidence from a Dividing Sudan." *The Journal of Politics* 76 (03): 649–665.

Berger, Jonah, Marc Meredith, and S Christian Wheeler. 2008a. "Contextual priming: Where people vote affects how they vote." *Proceedings of the National Academy of Sciences* 105 (26): 8846–8849.

Berger, Jonah, Marc Meredith, and S Christian Wheeler. 2008b. "Contextual Priming: Where People Vote Affects How they Vote." *Proceedings of the National Academy of Sciences* 105 (26): 8846–8849.

Bhavnani, Ravi, Karsten Donnay, Dan Miodownik, Maayan Mor, and Dirk Helbing. 2014. "Group segregation and urban violence." *American Journal of Political Science* 58 (1): 226–245.

Bobo, Lawrence, Camille L Zubrinsky, James H Johnson Jr, and Melvin L Oliver. 1994. "Public opinion before and after a spring of discontent." *The Los Angeles riots: Lessons for the urban future*: 103–133.

Botta, Federico, Helen Susannah Moat, and Tobias Preis. 2015. "Quantifying crowd size with mobile phone and Twitter data." *Royal Society open science* 2 (5): 150162.

Boustan, Leah Platt. 2013. Racial residential segregation in American cities. Technical report National Bureau of Economic Research.

Bratton, Michael, Ravi Bhavnani, and Tse-Hsin Chen. 2012. "Voting intentions in Africa: ethnic, economic or partisan?" *Commonwealth & Comparative Politics* 50 (1): 27–52.

Bratton, Michael, and Robert Mattes. 2003. "Support for economic reform? Popular attitudes in Southern Africa." *World Development* 31 (2): 303–323.

Burchfield, Marcy, Henry G Overman, Diego Puga, and Matthew A Turner. 2006. "Causes of sprawl: A portrait from space." *The Quarterly Journal of Economics*: 587–633.

California Governor's Commission on the Los Angeles Riots. 1963. *Violence in the City—An End or a Beginning? A Report by the Governor's Commission on the Los Angeles Riots*. Sacramento, CA: California Department of General Services.

Campbell, Angus, Philip E Converse, Warren E Miller, and Donald E Stokes. 1960. *The American Voter*. Wiley.

Campbell, Jennifer D, and Abraham Tesser. 1985. "Self-evaluation maintenance processes in relationships." In *Understanding personal relationships: An interdisciplinary approach*, ed. Steve Duck and Daniel Perlman. Thousand Oaks, CA: Sage Publications, Inc.

Carlson, Elizabeth. 2015. "Ethnic Voting and Accountability in Africa: A Choice Experiment in Uganda." *World Politics* 67 (02): 353–385.

Carpenter, Daniel, and Colin D Moore. 2014. "When Canvassers Became Activists: Antislavery Petitioning and the Political Mobilization of American Women." *American Political Science Review* 108 (03): 479–498.

Cavaillé, Charlotte, and Kris-Stella Trump. 2015. "The two facets of social policy preferences." *The Journal of Politics* 77 (1): 146–160.

CDC. 2016. "Health of Black or African American non-Hispanic Population."

Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity experiment." *The American Economic Review* 106 (4): 855–902.

Christopher, AJ. 2001. "Urban segregation in post-apartheid South Africa." *Urban Studies* 38 (3): 449–466.

Christopher, Anthony John. 1990. "Apartheid and urban segregation levels in South Africa." *Urban Studies* 27 (3): 421–440.

Christopher, Anthony John. 1994. *The atlas of apartheid*. Witwatersrand University Press Publications.

Clark, William AV, and Natasha Rivers. 2013. "Community choice in large cities: Selectivity and ethnic sorting across neighbourhoods." In *Understanding Neighbourhood Dynamics*. Springer.

CNN. 2013. "Los Angeles Riots Fast Facts." http://www.cnn.com/2013/09/18/us/los-angeles-riots-fast-facts/.

Cohen, Darryl T, GW Hatchard, and SG Wilson. 2015. "Population Trends in Incorporated Places: 2000 to 2013." *Population Estimates and Projections. Current Population Reports*.

Cohen, Geoffrey L. 2003. "Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs." *Journal of Personality and Social Psychology* 85 (5): 808 – 822.

Condon, Meghan, and Amber Wichowsky. 2016. "Inequality in the Social Mind: Social Comparison and Support for Redistribution." *Working Paper*.

Côté, Stéphane, Julian House, and Robb Willer. 2015. "High economic inequality leads higher-income individuals to be less generous." *Proceedings of the National Academy of Sciences*: 201511536.

Darley, John M, and Bibb Latane. 1968. "Bystander intervention in emergencies: diffusion of responsibility." *Journal of Personality and Social Psychology* 8 (4): 377–383.

Davies, RJ. 1967. "The South African urban hierarchy." *South African Geographical Journal* 49 (1): 9–20.

Davies, Ronald J. 1981. "The spatial formation of the South African city." *GeoJournal* 2 (2): 59–72.

Dawson, Michael C. 1995. *Behind the Mule: Race and Class in African-American Politics*. Princeton, NJ: Princeton University Press.

De Waal, Alex. 2005. "Who are the Darfurians? Arab and African Identities, Violence and External Engagement." *African Affairs* 104 (415): 181–205.

DeCelles, K.A., and M.I. Norton. 2016. "Physical and Situational Inequality on Airplanes Predict Air Rage." *Proceedings of the National Academy of Sciences* 113 (20): 5588–5591.

Dionne, Kim Yi, Darin DeWitt, Michael Stone, and Michael Suk-Young Chwe. 2014. "The May 1 Marchers in Los Angeles: Overcoming Conflicting Frames, Bilingual Women Connectors, English-Language Radio, and Newly Politicized Spanish Speakers." *Urban Affairs Review*: 1078087414534009.

Dunning, Thad, and Lauren Harrison. 2010. "Cross-cutting cleavages and ethnic voting: An experimental study of cousinage in Mali." *American Political Science Review* 104 (01): 21–39.

Edsall, Thomas Byrne, and Mary D. Edsall. 1992. *Chain Reaction: The Impact of Race, Rights, and Taxes on American Politics*. New York: Norton.

Edwards, Iain. 1994. "Cato Manor: cruel past, pivotal future." *Review of African Political Economy* 21 (61): 415–427.

Eifert, Benn, Edward Miguel, and Daniel N Posner. 2010. "Political competition and ethnic identification in Africa." *American Journal of Political Science* 54 (2): 494–510.

Emmons, David M. 1965. *The pedestrian count*. American Society of Planning Officials.

Enos, Ryan D. 2014. "The Causal Effect of Intergroup Contact on Exclusionary Attitudes." *Proceedings of the National Academy of Sciences of the United States of America* 111 (10): 3699-3704.

Enos, Ryan D. 2016. "What the demolition of public housing teaches us about the impact of racial threat on political behavior." *American Journal of Political Science* 60 (1): 123–142.

Enos, Ryan D, Anthony Fowler, and Lynn Vavreck. 2014. "Increasing Inequality: The Effect of GOTV Mobilization on the Composition of the Electorate." *Journal of Politics* 76 (1): 273–288.

Enos, Ryan D., and Christopher Celaya. 2015. "Segregation Directly Affects Human Perception and Intergroup Bias." American Political Science Association, annual meeting, Chicago.

Enos, Ryan D., and Noam Gidron. 2016. "Intergroup Behavioral Strategies as Contextually Determined: Experimental Evidence from Israel." *Journal of Politics* 78 (3): 851–867.

Erikson, Robert S., Michael MacKuen, and James A. Stimson. 2002. *The Macro Polity*. Cambridge Studies in Political Psychology and Public Opinion Cambridge: Cambridge University Press.

Evans, Ivan Thomas. 1997. *Bureaucracy and race: Native administration in South Africa*. Vol. 53 Univ of California Press.

Ferree, Karen E. 2006. "Explaining South Africa's racial census." *Journal of Politics* 68 (4): 803–815.

Ferree, Karen E. 2010. *Framing the race in South Africa: the political origins of racial census elections*. Cambridge University Press.

Festinger, Leon. 1954. "A theory of social comparison processes." *Human relations* 7 (2): 117–140.

Fiske, Susan T. 2000. "Stereotyping, Prejudice, and Discrimination at the Seam Between the Centuries: Evolution, Culture, Mind, and Brain." *European Journal of Social Psychology* 30 (3): 299–322.

Flint, Colin. 1996. "Whither the individual, whither the context?" *Political Geography* 15 (2): 147–151.

Fogelson, Robert M, and Robert Bernard Hill. 1968. *Who Riots?: A Study of Participation in the 1967 Riots*. US Government Printing Office.

Frank, Lawrence D, and Peter O Engelke. 2001. "The built environment and human activity patterns: exploring the impacts of urban form on public health." *Journal of Planning Literature* 16 (2): 202–218.

Friedman, Steven. 2004. "Why we vote: the issue of identity." *Election Synopsis* 1 (2): 2–4.

Fruin, John J. 1971. Pedestrian planning and design. Technical report.

Gaertner, Lowell, Constantine Sedikides, and Kenneth Graetz. 1999. "In search of self-definition: motivational primacy of the individual self, motivational primacy of the collective self, or contextual primacy?" *Journal of Personality and Social Psychology* 76 (1): 5.

Gay, Claudine. 2012. "Moving to Opportunity: The Political Effects of a Housing Mobility Experiment." *Urban Affairs Review* 48 (2): 147–179.

Gibson, James L, and Amanda Gouws. 2005. *Overcoming intolerance in South Africa: Experiments in democratic persuasion.* Cambridge University Press.

Gibson, James L, and Christopher Claassen. 2010. "Racial reconciliation in South Africa: Interracial contact and changes over time." *Journal of Social Issues* 66 (2): 255–272.

Gilens, Martin. 1996. "'Race coding' and white opposition to welfare." *American Political Science Review* 90 (03): 593–604.

Gilens, Martin. 1999. *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy.* Chicago: University of Chicago Press.

Gilens, Martin. 2012. *Affluence and Influence: Economic Inequality and Political Power in America.* Princeton, NJ: Princeton University Press.

Goffman, Erving. 1971. "Relations in public: Microstudies of the social order." *London: Allen Lane.*

Green, Donald, Bradley Palmquist, and Eric Schickler. 2002. *Partisan Hearts and Minds.* New Haven: Yale University Press.

Greenberg, Jeff, Tom Pyszczynski, and Sheldon Solomon. 1986. "The Causes and Consequences of a Need for Self-esteem: A Terror Management Theory." In *Public self and private self.* Springer.

Habyarimana, James, Macartan Humphreys, Daniel N. Posner, and Jeremy M. Weinstein. 2009. *Coethnicity: Diversity and the Dilemmas of Collective Action.* New York: Russell Sage.

Han, Hahrie. 2016. "The organizational roots of political activism: Field experiments on creating a relational context." *American Political Science Review* 110 (02): 296–307.

Hayes, Bernadette C, and Ian McAllister. 2001. "Sowing Dragon's Teeth: Public Support for Political Violence and Paramilitarism in Northern Ireland." *Political Studies* 49 (5): 901–922.

Hopkins, Daniel J. 2009. "The Diversity Discount: When Increasing Ethnic and Racial Diversity Prevents Tax Increases." *The Journal of Politics* 71 (01): 160–177.

Hopkins, Daniel J. 2010. "Politicized Places: Explaining Where and When Immigrants Provoke Local Opposition." *American Political Science Review* 104 (1): 40–60.

Horowitz, Donald L. 1985. *Ethnic Groups in Conflict.* University of California Press.

Huntington, Samuel P. 1968. *Political Order in Changing Societies.* New Haven, CT: Yale University Press.

Ichino, Nahomi, and Noah L Nathan. 2013. “Crossing the line: Local ethnic geography and voting in Ghana.” *American Political Science Review* 107 (02): 344–361.

Ihlanfeldt, Keith R, and Benjamin Scafidi. 2002. “Black self-segregation as a cause of housing segregation: Evidence from the multi-city study of urban inequality.” *Journal of Urban Economics* 51 (2): 366–390.

Ishiyama, John. 2012. “Explaining ethnic bloc voting in Africa.” *Democratization* 19 (4): 761–788.

Iyengar, Shanto. 1990. “Framing responsibility for political issues: The case of poverty.” *Political Behavior* 12 (1): 19–40.

Jacobs, Jane. 1961. *The Death and Life of Great American Cities*. New York: Random House.

John, Peter, Sarah Cotterill, Liz Richardson, Alice Moseley, Gerry Stoker, Graham Smith, Corinne Wales, Hanhua Liu, and Hisako Nomura. 2013. *Nudge, nudge, think, think: Experimenting with ways to change civic behaviour*. A&C Black.

Johnson, Richard William, and Lawrence Schlemmer. 1996. *Launching Democracy in South Africa: the first open election, April 1994*. Yale University Press.

Kasara, Kimuli. 2013. “Separate and suspicious: Local social and political context and ethnic tolerance in Kenya.” *The Journal of Politics* 75 (04): 921–936.

Kasara, Kimuli. 2014. “Does Local Ethnic Segregation Lead to Violence?: Evidence from Kenya.”.

Katz, Lawrence F, Jeffrey R Kling, and Jeffrey B Liebman. 2001. “Moving to opportunity in Boston: Early results of a randomized mobility experiment.” *The Quarterly Journal of Economics* 116 (2): 607–654.

Key, V.O. 1949. *Southern Politics in State and Nation*. New York: Knopf.

Khan, Sultan, and Brij Maharaj. 1998. “Restructuring the apartheid city.” In *Urban Forum*. Vol. 9 Springer pp. 197–223.

Kim-Goh, Mikyong, Chong Suh, Dudley David Blake, and Bruce Hiley-Young. 1995. “Psychological impact of the Los Angeles riots on Korean-American victims: Implications for treatment.” *American Journal of Orthopsychiatry* 65 (1): 138.

Kinder, Donald R, and David O Sears. 1981. “Prejudice and politics: Symbolic racism versus racial threats to the good life.” *Journal of Personality and Social Psychology* 40 (3): 414.

King, Gary. 1996. “Why Context Should Not Count.” *Political Geography* 15: 159–164.

King, Gary. 1997. *A Solution to the Ecological Inference Problem: Reconstructing Individual Behavior from Aggregate Data*. Princeton, NJ: Princeton University Press.

Kling, Jeffery R., Jeffry B. Liebman, and Lawrence F. Katz. 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica* 75 (1): 83–119.

Kuper, Leo, Hilstan Lett Watts, and Ronald Davies. 1958. *Durban: a study in racial ecology.* J. Cape.

Latané, Bibb. 1981. "The Psychology of Social Impact." *American Psychologist* 36 (4): 343–356.

Lee, Yu-Hao, and Gary Hsieh. 2013. "Does slacktivism hurt activism?" In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM pp. 811–820.

Lenz, Gabriel S. 2013. *Follow the Leader?: How Voters Respond to Politicians' Policies and Performance*. Chicago: University of Chicago Press.

Lerman, Amy E, and Vesla Weaver. 2014. "Staying out of sight? Concentrated policing and local political action." *The ANNALS of the American Academy of Political and Social Science* 651 (1): 202–219.

Lerner, Melvin J. 1980. *The belief in a just world*. Springer.

Levine, Jeremy R, and Carl Gershenson. 2014. "From Political to Material Inequality: Race, Immigration, and Requests for Public Goods." In *Sociological Forum*. Vol. 29 Wiley Online Library pp. 607–627.

Leyden, Kevin M. 2003. "Social capital and the built environment: the importance of walkable neighborhoods." *American Journal of Public Health* 93 (9): 1546–1551.

Lieberman, Evan S, and Gwyneth H McClendon. 2013. "The ethnicity–policy preference link in sub-Saharan Africa." *Comparative Political Studies* 46 (5): 574–602.

Lodge, Tom. 1999. *Consolidating Democracy: South Africa's Second Popular Election*. Witwatersrand University Press.

Louw, P Eric. 2000. "South Africa's second post-apartheid elections: A reaffirmation of racial politics but a loosening of ethnic bonds." *Australian Journal of International Affairs* 54 (2): 217–238.

Løvås, Gunnar G. 1994. "Modeling and simulation of pedestrian traffic flow." *Transportation Research Part B: Methodological* 28 (6): 429–443.

Lupu, Noam, and Jonas Pontusson. 2011. "The structure of inequality and the politics of redistribution." *American Political Science Review* 105 (02): 316–336.

Manza, Jeff, and Christopher Uggen. 2006. *Locked Out: Felon Disenfranchisement and American Democracy*. Oxford University Press, USA.

Massey, Douglas S. 1993. *American Apartheid: Segregation and the Making of the Underclass*. Harvard University Press.

Massey, Douglas S, Gretchen A Condran, and Nancy A Denton. 1987. "The effect of residential segregation on black social and economic well-being." *Social Forces* 66 (1): 29–56.

Massey, Douglas S., and Nancy A. Denton. 1988. "The Dimensions of Residential Segregation." *Social Forces* 67.

Mattes, Robert B. 1995. *The Election Book: Judgement and choice in South Africa's 1994 election*. IDASA, Public Information Centre.

Mattes, Robert, and Jessica Piombo. 2001. "Opposition parties and the voters in South Africa's general election of 1999." *Democratization* 8 (3): 101–128.

Meltzer, Allan H, and Scott F Richard. 1981. "A rational theory of the size of government." *The Journal of Political Economy*: 914–927.

Meredith, Marc. 2009. "Persistence in Political Participation." *Quarterly Journal of Political Science* 4 (3): 187–209.

Merton, Robert K. 1957. *Social Theory and Social Structure*. Free Press.

Miller, Carol T. 1982. "The role of performance-related similarity in social comparison of abilities: A test of the related attributes hypothesis." *Journal of Experimental Social Psychology* 18 (6): 513–523.

Miller, Carol T. 1984. "Self-schemas, gender, and social comparison: A clarification of the related attributes hypothesis." *Journal of Personality and Social Psychology* 46 (6): 1222.

Navarrete, Carlos David, and Daniel MT Fessler. 2005. "Normative Bias and Adaptive Challenges: A Relational Approach to Coalitional Psychology and a Critique of Terror Management Theory." *Evolutionary Psychology*.

Newman, Benjamin J, Christopher D Johnston, and Patrick L Lown. 2015. "False Consciousness or Class Awareness? Local Income Inequality, Personal Economic Position, and Belief in American Meritocracy." *American Journal of Political Science* 59 (2): 326–340.

Nishi, Akihiro, Hirokazu Shirado, David G Rand, and Nicholas A Christakis. 2015. "Inequality and visibility of wealth in experimental social networks." *Nature*.

Nunn, Nathan, and Diego Puga. 2012. "Ruggedness: The blessing of bad geography in Africa." *Review of Economics and Statistics* 94 (1): 20–36.

O'Brien, Daniel Tumminelli, Dietmar Offenhuber, Jessica Baldwin-Philippi, Melissa Sands, and Eric Gordon. 2016. "Uncharted Territoriality in Coproduction: The Motivations for 311 Reporting." *Journal of Public Administration Research and Theory* 27.

Oliver, J Eric. 1999. "The effects of metropolitan economic segregation on local civic participation." *American Journal of Political Science*: 186–212.

Oliver, J Eric. 2001. *Democracy in Suburbia*. Princeton University Press.

O'Loughlin, John. 2000. "Can King's Ecological Inference Method Answer a Social Scientific Puzzle: Who Voted for the Nazi Party in Weimar Germany?" *Annals of the Association of American Geographers*: 592–601.

Olzak, Susan, Suzanne Shanahan, and Elizabeth H McEneaney. 1996. "Poverty, Segregation, and Race Riots: 1960 to 1993." *American Sociological Review*: 590–613.

O'Sullivan, David, and David Unwin. 2014. *Geographic Information Analysis*. John Wiley & Sons.

Parry, Janine A, Daniel A Smith, and Shayne Henry. 2012. "The impact of petition signing on voter turnout." *Political Behavior* 34 (1): 117–136.

Peffley, Mark, Jon Hurwitz, and Paul M Sniderman. 1997. "Racial stereotypes and whites' political views of blacks in the context of welfare and crime." *American Journal of Political Science*: 30–60.

Perdikaki, Olga, Saravanan Kesavan, and Jayashankar M Swaminathan. 2012. "Effect of traffic on sales and conversion rates of retail stores." *Manufacturing & Service Operations Management* 14 (1): 145–162.

Perry, F. Noel, Sarah Henry, Marcia E. Perry, and Stephanie Leonard. 2016. "California Migration: A Comparative Perspective."

Pettigrew, Stephen. forthcoming. "The Race Gap in Precinct Wait Times: Why Minority Precincts are Underserved by Local Election Officials." *Political Science Quarterly*.

Piketty, T, and E Saez. 2013. "Income Inequality in the United States, 1913-1998 (Tables and figures updated to 2012) Retrieved 24 January, 2014."

Piketty, Thomas. 1995. "Social mobility and redistributive politics." *The Quarterly Journal of Economics*: 551–584.

Piketty, Thomas. 1998. "Self-fulfilling beliefs about social status." *Journal of Public Economics* 70 (1): 115–132.

Posner, Daniel N. 2004. "The political salience of cultural difference: Why Chewas and Tumbukas are allies in Zambia and adversaries in Malawi." *American Political Science Review* 98 (4): 529–545.

Posner, Daniel N, and David J Simon. 2002. "Economic Conditions and Incumbent Support in Africa's New Democracies: Evidence from Zambia." *Comparative Political Studies* 35 (3): 313–336.

Pushkarev, Boris. 1975. "Urban space for pedestrians.".

Reicher, Stephen D, Anne Templeton, Fergus Neville, Lucienne Ferrari, and John Drury. 2016. "Core Disgust is Attenuated by Ingroup Relations." *Proceedings of the National Academy of Sciences*: 201517027.

Ribeiro Jr., P.J., and P.J. Diggle. 2001. "geoR: a package for geostatistical analysis." *R-NEWS* 1 (2): 15–18.

Rieder, Jonathan. 1985. *Canarsie: The Jews and Italians of Brooklyn Against Liberalism.* Cambridge, MA: Harvard University Press.

Riley, Shawn J, Stephen D DeGloria, and Robert Elliot. 1999. "A terrain ruggedness index that quantifies topographic heterogeneity." *Intermountain Journal of Sciences* 5 (1-4): 23–27.

Roberts, Kevin WS. 1977. "Voting over income tax schedules." *Journal of Public Economics* 8 (3): 329–340.

Robinson, WS. 1950. "Ecological Correlations and the Behavior of Individuals." *American Sociological Review* 15: 351–357.

Romer, Thomas. 1975. "Individual welfare, majority voting, and the properties of a linear income tax." *Journal of Public Economics* 4 (2): 163–185.

Rosenthal, Robert, and Kermit L Fode. 1963. "The effect of experimenter bias on the performance of the albino rat." *Behavioral Science* 8 (3): 183–189.

Rugh, Jacob S, and Jessica Trounstine. 2011. "The Provision of Local Public Goods in Diverse Communities: Analyzing Municipal Bond Elections." *The Journal of Politics* 73 (04): 1038–1050.

Sands, Melissa. 2017a. "Replication Data for: Enos, Kaufman & Sands, Can violent protest change local policy support? Evidence from the aftermath of the 1992 Los Angeles riot." Harvard Dataverse. 10.7910/DVN/2MRBJL. http://dx.doi.org/10.7910/DVN/2MRBJL.

Sands, Melissa. 2017b. "Replication Data for: Exposure to inequality affects support for redistribution." Harvard Dataverse. 10.7910/DVN/E6BHOR. http://dx.doi.org/10.7910/DVN/E6BHOR.

Sands, Melissa. 2017c. "Replication Data for: 'Eyes' on the Street: What public camera feed data can teach us about civic and political behavior." Harvard Dataverse. 10.7910/DVN/HHWL1K. http://dx.doi.org/10.7910/DVN/HHWL1K.

Sands, Melissa. 2017d. "Replication Data for: Sands & de Kadt, Segregation drives racial voting: New evidence from South Africa." Harvard Dataverse. 10.7910/DVN/HF0UYN. http://dx.doi.org/10.7910/DVN/HF0UYN.

Sands, Melissa L. 2017e. "Exposure to inequality affects support for redistribution." *Proceedings of the National Academy of Sciences* 114 (4): 663–668.

Schelling, Thomas C. 1971a. "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1 (2): 143–186.

Schelling, Thomas C. 1971b. *Micromotives and Macrobehavior*. New York: Norton.

Schultz, Duane P. 1969. "The human subject in psychological research." *Psychological Bulletin* 72 (3): 214.

Sears, David O., and John B. McConahay. 1973. *The Politics of Violence: the New Urban Blacks and the Watts Riot*. Boston: Houghton Mifflin.

Sidanius, Jim, and Felicia Pratto. 2001. *Social Dominance: An Intergroup Theory of Social Hierarchy and Oppression*. New York: Cambridge University Press.

Simmel, Georg. 1978. "The philosophy of money, trans. Tom Bottomore and David Frisby." *London: Routledge & Kegan Paul* 236: 232.

Smith, Heather J, and Tom R Tyler. 1996. "Justice and power: When will justice concerns encourage the advantaged to support policies which redistribute economic resources and the disadvantaged to willingly obey the law?" *European Journal of Social Psychology* 26 (2): 171–200.

Sparks, Allister Haddon. 1990. *The Mind of South Africa*. Vol. 155 Heinemann London.

Spencer, Douglas M, and Zachary S Markovits. 2010. "Long lines at polling stations? Observations from an election day field study." *Election Law Journal* 9 (1): 3–17.

Telles, Edward E. 1992. "Residential segregation by skin color in Brazil." *American Sociological Review*: 186–197.

Tierney, Kathleen. 1994. "Property Damage and Violence: A Collective Behavior Analysis." *The Los Angeles Riots: Lessons for the Urban Future*: 149–173.

Toner, Robin. 1992. "After the Riots: Los Angeles Riots Are a Warning, Americans Fear." *New York Times*.

Trounstine, Jessica. 2015. "Segregation and Inequality in Public Goods." *American Journal of Political Science* 60.

Trump, Kris-Stella. 2017. "Income Inequality Influences Perceptions of Legitimate Income Differences." *British Journal of Political Science*: 1–24.

United States National Advisory Commission on Civil Disorders. 1968. *Report of the National Advisory Commission on Civil Disorders*. Washington, D.C.: US Government Publishing Office.

University of California, Los Angeles. Institute for Social Science Research. 2011. "Los Angeles County Social Survey, 1992 (LACSS)."

Van Kempen, Ronald, and A şule Özüekren. 1998. "Ethnic segregation in cities: new forms and explanations in a dynamic world." *Urban Studies* 35 (10): 1631–1656.

Verba, Sidney, Kay Lehman Schlozman, Henry E Brady, and Henry E Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics*. Vol. 4 Cambridge Univ Press.

Wang, Jingbo, Wai Wan Tsang, and George Marsaglia. 2003. "Evaluating Kolmogorov's distribution." *Journal of Statistical Software* 8 (18).

Wasow, Omar. 2016. "Do Protest Tactics Matter? Effects of the 1960s Black Insurgency on White Attitudes and Voting Behavior.". Working Paper: Princeton University.

Weghorst, Keith R, and Staffan I Lindberg. 2013. "What Drives the Swing Voter in Africa?" *American Journal of Political Science* 57 (3): 717–734.

Western, Bruce. 2006. *Punishment and Inequality in America*. Russell Sage Foundation.

Williamson, John B. 1974. "Beliefs about the motivation of the poor and attitudes toward poverty policy." *Social Problems* 21 (5): 634–648.

Wolff, Michael. 1973. "Notes on the behaviour of pedestrians." *People in Places: The Sociology of the Familiar*: 35–48.

Wood, Joanne V. 1989. "Theory and research concerning social comparisons of personal attributes." *Psychological Bulletin* 106 (2): 231.

Worden, Nigel. 2011. *The making of modern South Africa: conquest, apartheid, democracy*. John Wiley & Sons.

Young, Crawford. 1976. *The Politics of Cultural Pluralism*. University of Wisconsin Press.

Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.

Zorlu, Aslan, and Jan Latten. 2009. "Ethnic sorting in the Netherlands." *Urban Studies* 46 (9): 1899–1923.