What tearing down public housing projects teaches us about the effect of racial threat on political participation

Ryan D. Enos[1]

[1]Department of Political Science, University of California, Los Angeles; renos@ucla.edu.

Abstract

How does the context in which a person lives affect their political behavior? I exploit a natural experiment in which voters' demographic context was exogenously changed. Between 2000 and 2004, the reconstruction of public housing in Chicago caused the displacement of over 25,000 African Americans, many of whom had previously lived in close proximity to white voters. The removal was a largely systematic process, exogenous to the neighborhood of the public housing and even to the city. After the removal of their African American neighbors, the turnout of white voters dropped by over 5 percentage points. Additionally, the size of the effect decreases as the distance of the voters from the demolished projects increases. I develop a theory of spatial impact to explain how group spatial structure conditions the effect of context on behavior. The effect of the demolition of the housing projects is consistent with this theory of spatial impact and is evidence for Key's (1949) theory of racial threat.

One of the most significant demographic changes in United States history was the migration of African-Americans from the South to northern and western cities in the mid-Twentieth century. Scholars of the 1960's have claimed that the political reaction of urban whites to the influx of African Americans was palpable: previously apolitical individuals became politically activated (Rieder 1985, Edsall & Edsall 1992). For example, Edsall & Edsall (1992) documented racially liberal Paul Douglas' losses in white Chicago wards surrounding the expanding Black ghetto. The segregationist George Wallace was relatively successful in these same wards. The implication of these findings is that the individual behavior of these white voters was conditioned by the context in which they lived. But how do social scientists know if the context of surroundings really does condition behavior? Recently there has been renewed interest in studying the effects of context from the perspective of both social groups and geography. Research on social networks has indicated that the networks of which a person is part affect attitude formation (Huckfeldt & Sprague 1987) and political participation (Mutz 2002). The characteristic that receives the most attention in the study of American politics is geographically-based racial context. Voters are thought to behave differently depending on the racial demographics of a community (Oliver & Mendelberg 2000, Gay 2004).

However, the study of individual behavior and geographic context is difficult for two reasons. First, inference can be dependent on the choice of unit of measurement. Researchers will sometimes choose geographic units of measurement out of convenience because theories of context are sometimes silent with respect to scale. More often, data are only available for administrative units, such as Census Tracts. These units may have no natural social meaning and the correlation of an areal unit with individual behavior may change with the unit chosen by the researcher. For example, research on racial homogeneity and voter behavior has been modeled as being dependent on the racial composition of the state (Leighley & Nagler 1992), county (Giles & Buckner 1993), zip code (Leighley & Vedlitz 1999), and Census Tract (Putnam 2007). While these geographies can be important, there is little reason to suppose that living in a diverse state has the same effect as living in a diverse zip code.

As such, researchers trying to measure the same effect can come to very different answers depending on their choice of unit. Second, and perhaps most seriously, self-selection makes it difficult to establish a causal connection between context and behavior. Individuals almost always have some degree of autonomy about where they choose to be, so it is difficult to separate the causal effect of context from other variables that lead the person to that context in the first place. For example, individuals in racially diverse neighborhoods have more liberal racial attitudes (Farley, Steeh, Krysan, Jackson, & Reeves 1994), but were these attitudes what caused them to move to these neighborhoods or did they develop these attitudes because of the neighborhood? Even when individual-level characteristics are controlled for, it remains unclear if attributes of a person's neighborhood directly affect behavior or simply reflect unmeasured characteristics of the person. In order to discover the influence of geographic context on individual behavior, both of these challenges must be addressed.

In this paper, I exploit a natural experiment that is not tied to pre-existing areal units and in which context is changed abruptly and exogenously. The demolition of twelve large public-housing projects in Chicago, starting around 2001, removed roughly 25,000 people from Chicago neighborhoods in which they had previously lived. Notably, nearly all of these families were African American. Several of these demolished housing projects were in close proximity to predominately white neighborhoods. The demolition of these projects precipitated a large-scale demographic change to the surrounding neighborhoods. Because the decision to demolish these projects was outside the control of those who lived near (or in) the projects and because I test for a change in behavior before substantial resorting could take place, I am able to separate the behavioral effect of the demographic change from other preexisting influences and from self-selection in or out of the neighborhood. A comparison of voting rates before and after the demolition of the projects provides a test of the influence of the housing project and its African American residents (context) on the turnout of nearby white voters (political behavior). When African Americans are removed from the neighborhood, how do white voters respond?

As described in Section 4, after the demolition of the projects, white voters' turnout dropped by more than five percentage points for those voters living at the closest distances from the projects. The turnout of African Americans living nearby did not change. This result is maintained even when a number of alternative tests are considered.

The rest of the paper is organized as follows: Section 1 describes group threat, which is a theory of contextual influence that I am able to test with this natural experiment. In the same section, I describe my theory of spatial impact and how I am able to test that as well. Section 2 gives a brief overview of public housing in Chicago. Section 3 describes the research design for the study and how I estimate the race of voters from surnames and Census demographics. The final two sections present the results of the analysis and discuss the wider implications of the findings.

1 Group threat and the study of context

In the American politics literature, the most commonly studied topic of residential context is the influence of racial demographics. Key's (1949) theory of racial treat set the stage for a long line of research. Key found that, at the county level in the United States South, white voter turnout was correlated with the number of African Americans in the county. Key claimed that whites were threatened by the presence of African Americans and, therefore, were more politically motivated. Blumer (1958) described a more general theory of group threat in which members of a group react negatively when they perceive that their place in the social hierarchy is threatened by an inferior group (see also Blalock (1967) and Bobo & Hutchings (1996)). Since Key's finding, subsequent research has questioned both whether the theory is portable outside the post-war American South and whether racial threat exists at all, or if the apparent effect of racial composition is confounded with other factors (Voss 1996).

Subsequent research on racial threat has yielded mixed results. Studies that have found no evidence for group threat have included Brewer & Miller (1988), Ellison & Powers (1994),

Sigelman & Welch (1993), Welch, Sigelman, Bledsoe, & Combs (2001), and Fitzpatrick & Hwang (1992). Leighley & Vedlitz (1999) even find that white turnout is *negatively* correlated with group threat, while group threat has different effects for different ethnic groups. And Voss (1996), in a critique of Giles & Buckner (1993), shows that detecting group threat is dependent on the choice of areal unit. However, group threat is supported in the studies by Matthews & Prothro (1963), Carsey (1995), Glaser (1994), Fossett & Kiecolt (1989), Quillian (1995), Taylor (1998), and Wright (1977), among others.

That the question of whether group threat influences behavior remains largely unsettled speaks to the difficulties associated with the study of context in general and group threat in particular. I argue that part of the difficulty comes from an incomplete theoretical framework of the effect of context.

1.1 Spatial Impact

All of the above mentioned studies, and nearly all studies of demographic context, model context as dependent on a single attribute: group size. For example, what proportion of a county is African American? However, group size alone does not fully capture the context that conditions racial threat. Rather, I argue that geographic context, and racial threat in particular, is a function of three dimensions: relative group size, group proximity, and spatial concentration. To understand the threat posed by a group, we must know more than just the size of a group. We must also consider how groups relate spatially: how close they are (proximity) and how spatially segregated the groups are (concentration). These three dimensions together - size, proximity, and concentration - interact to produce what I call spatial impact. Like theories of group threat, I argue that spatial impact can result in increased threat from the outgroup that can lead to increased individual political activity.

Like theories of group threat, the mechanism driving spatial impact is a psychological reaction to a perceived threat. Group competition arises easily when group distinctions are made within a population, even if the distinction is completely arbitrary (Tajfel, Billig,

Bundy, & Flament 1971). I am proposing that if two outgroups are of equal size, the one that is spatially closer, is perceived as more of a threat. And if two groups are of equal size and equidistant, the group that is more concentrated and spatially distinct from other groups is also more of a threat than one that is spread out and lacking strength. This follows from the psychological structure of the impact of social stimuli generally. Latané (1981) shows that any social stimulus has an impact that is a function of size, strength, and immediacy (see also Latané & Wolf (1981)). These dimensions are analogous to the dimensions that affect the impact of groups in space. The size of the social stimulus is analogous to the size of the group in space; the immediacy of the social stimulus is analogous to the spatial proximity of a group; and the strength of social stimulus becomes the spatial concentration.[1]

High levels of size, concentration, and proximity describes the neighborhoods surrounding public housing projects in Chicago that are the focus of this study. As such, the demolition of these projects provides an excellent test of effect of spatial impact by measuring what happens after the projects were removed.

2 Chicago Public Housing

Since 1999, the Chicago Housing Authority (CHA) has relocated thousands of families through the process of destroying and reconfiguring its massive system of public housing in the city.

Prior to 2000, when widespread demolition begin, the CHA was the second largest public housing agency in the United States. The agency controlled over 2800 properties.[2] The great majority of the properties were considered "scattered-site", meaning they were units consisting of a single building, unattached to other public housing. However, 83 properties were large, multi-building properties that, collectively, housed tens-of-thousands of families. All of the families in this housing were low-income and the overwhelming majority were

[1] I have also argued that increased spatial concentration leads to increased efficiency of communication within group and this can lead to heightened political activity.

[2] Sixty-nine of these units were senior housing.

African-American. Of the demolished projects for which data is available, the average racial composition was 99.7% Black. The great majority of housing was located on the city's South and West sides. And most of the housing was located in neighborhoods that were predominately Latino or African American. Additionally, because Chicago is one of the most segregated cities in the United States (Massey & Denton 1993), most of these neighborhoods were almost exclusively African American or Latino. However, there were exceptions so that projects were located close to predominately white neighborhoods. I use this proximity to white neighborhoods to study the spatial impact of the population in the housing projects on the white voters surrounding the projects.

Low income public housing was designated by the CHA, in their 1997 "Plan for Transformation" as requiring demolition or not. The guidelines for requiring demolition were set by the United States Department of Housing and Urban Development (HUD), in Washington, D.C.. The process was, more or less, decided exogenously to the particular neighborhood. Rules deciding demolition were based on measures of size of facilities and levels of decay that were outside of the control of residents (Authority 2000).[3] Some housing was recon-

[3]Residents of public housing who were relocated were moved to what is known as "Section 8" housing, which is where residents are provided with vouchers to help pay for housing in private facilities. In demolishing the projects, HUD and the CHA were following the law from Section 202 of the Omnibus Consolidated Rescissions and Appropriations Act of 1996. An internal HUD audit describes the relevant part of Section 202 as follows:

> The public housing developments that were subject to Section 202 must have been on the same or contiguous sites, must have been more expensive than Section 8 tenant-based assistance, could not have been revitalized through reasonable programs, and must have had 300 or more dwelling units and had a vacancy rate of at least 10 percent for dwelling units not in funded on-schedule modernization programs.
>
> Section 202 requires that a public housing development's cost of operations be analyzed to determine whether it is more expensive to renovate and operate the low-income housing units than it is to provide Section 8 assistance to current residents and operate the low-income housing units than it is to provide Section 8 assistance to current residents and relocate them to other available developments. For a public housing development for which revitalization is deemed the more expensive option, additional analysis is undertaken to assess its long-term viability if revitalized. Under current definitions, viability includes achieving structural/system soundness and full occupancy. These factors, along with an analysis of market support, must be considered to establish the appropriate long-range plan for the use or reuse of a public housing development. Low-income housing units that do not pass the tests involved in this analysis and that cannot be revitalized through reasonable programs are required by Section 202 to be removed from HUD's public housing inventory (Wolfe 2006).

stituted and some was entirely demolished. While almost all the reconstruction required the displacement of residents, the units requiring demolition were overwhelmingly the large high-rise, multi-building complexes that had become notorious for poor living conditions in American cities. The destruction of these units changed the demography and density of the neighborhoods in which they previously existed.[4] There were twelve projects in the CHA that were scheduled for demolition and completely or partially demolished between 2000 and 2004, leading to the displacement of over 25,000 people.[5]

The geography and timing of the demolition of the projects allow for a straightforward test of the impact of the African American outgroup on white voters. All that needs to be done is to compare the difference in turnout, relative to the rest of the city, of residents that lived close to the projects in 2000 and 2004. If, relative to other white residents, the turnout of whites close to the demolished projects is lower in 2004 than 2000, this is evidence of change due to the exogenous shock to their demographic context.

3 Design

The tearing down of the large-scale public housing projects can be thought of as a quasi-experiment. The treatment is the demolition of the housing projects and the outcome is the level of white political participation. In order to undertake this test, I obtained the

[4] Actually, the demographic change was probably more dramatic than even the 25,000 person displacement mentioned previously. This is because CHA resident counts do not include the thousands of homeless squatters and residents illegally living with legal residents. Some estimates have put the proportion of occupied units that were illegally occupied at as high as 50% (see Kotlowitz (1992) and Cunningham, Popkin, Eiseman, & Ferryman (2005)).

[5] These were:

Frances Cabrini Extension North
Ida B. Wells Homes and Wells Extension/ Madden Park Homes/ Clarence Darrow Homes
Frances Cabrini Extension South
William Green Homes
Rockwell Gardens
Stateway Gardens
Robert Taylor Homes (A and B)
Washington Park Homes (high-rises)
Henry Horner Homes
ABLA (Jane Addams Homes, Robert H. Brooks Homes and Brooks Extension, Loomis Courts, and the Grace Abbott Apartments)

2004 Illinois voter file and augmented it with Census demographics at the Census Block and Block Group level. I then geocoded the residences of the approximately 1.2 million voters in Chicago and determined their distance from each of the demolished public housing projects using a Geographic Information System. I also identified the race of each voter using a Bayesian process that estimates a voter's race based on their name and location. Because I am able to identify the race and location of each individual voter, instead of being tied to administrative units, I can test for an effect of the treatment using a variety of definitions of the relevant context. This means that I do not have to confine this analysis to a certain scale of measurement.

The difficulty in choosing scale of measurement in studies with geographic units is widely recognized by geographers. A related problem, also widely recognized by geographers is the Modifiable Areal Unit Problem (MAUP). MAUP refers not to the choice of scale but to the arbitrary boundary of some areal units, such as Census Tracts. Both MAUP and the problem of scale lead to the same potential difficulty with inference: simply that when using aggregate data to measure an underlying population parameter, the choice of the aggregate areal unit can matter as much as the underlying variation in population characteristics. For example, had Key (1949) chosen to undertake his study at the city level, rather than the county level, would his findings have been different? Or if Key had chosen the Census Tract as the unit of analysis, the essentially arbitrary boundaries of the Census Tracts may drive the results. This difficulty is generally present for any study of context that uses pre-defined boundaries, especially since many administrative boundaries often have no social meaning. Because I individually geocode each voter, I am able to test the effect of the destruction of the projects at any distance. Had I not individually geocoded voters, I would have to conduct the analysis using a unit like Census Tracts and, thus, the analysis would be susceptible to the Modifiable Areal Unit Problem.

Formally, the theory of spatial impact implies the following central hypotheses that I will test in this paper:

H1 group threat: After the demolition of the projects, if white voters close to the projects vote less, relative to the rest of the city, this is evidence for group threat.

H2 proximity: This design allows me to test an implication of the theory of spatial impact: that impact should vary with proximity. The effect of the treatment should decline as the white voters are further away from the project.

In addition to proximity, spatial impact is also dependent on group size and concentration. The situation of public housing in Chicago prior to its demolition is an example of high levels of both of these variables. Although I am not able to manipulate the levels of size and concentration, the spatial impact of the public housing before demolition should be very high. The housing projects consisted of a large outgroup, African Americans, that were extremely spatially-concentrated into imposing high-rise buildings. Additionally, many projects were very close, within just a few hundred meters of predominately white neighborhoods. If spatial impact is not demonstrated in this situation, then it is unlikely that the theory holds elsewhere.

3.1 Identifying the race of voters

A key variable in this analysis is race. I must differentiate between white, Black, and other voters. However, a voters' race is not identified on the public voting file in Illinois, which I used to conduct the analysis, so I develop a method to estimate the race of voters. Most previous studies that attempt to estimate the race of voters have used either Census demographics from where a voter lives or have used lists identifying how often a voters' surname comes from a particular racial group. Using only names has the shortcoming that there is considerable overlap between the names shared by some racial groups: whites and African Americans, for example, often have similar surnames. Using Census information has the disadvantage that the same estimates must be applied to everyone in the Census

geography. Instead, I use a method to estimate the race/ethnicity of voters that combines Census demographics with name frequencies. I apply Bayes' Rule to update priors on both surname frequency by race and Census block demographics to estimate the probability that the voter is white, Black, or any of the Census racial and ethnic categories. I describe this process in more detail in the Appendix.

Because of the hyper-segregated nature of Chicago, like most large United States cities (Massey & Denton 1993), I am able to make very certain predictions about the racial identity of many of the voters. This is because so many Census Blocks are either overwhelmingly Black or overwhelmingly white.

As an example, I randomly selected five voters from my sample. They are listed in Table 1. The probability that each of these voters is either white or Black is quite high. I include here only the probabilities for Black and white, although the frequencies of all five of the Census designated racial groups and Hispanic are used to estimate the probabilities. The names selected are not atypical for a segregated city, such as Chicago, and show how updating on the name and the Census demographic information gives more certain probabilities than if location or name were used alone.

Looking at an example from the Table 1 can help to demonstrate how this method improves estimates of race/ethnicity for individuals where using either just census demographics or just their name might be misleading. For example, the first voter in Table 1 has the last name Collins. Using only the name Collins, this voter would likely be white, with over 73% of people in the United States with the name Collins being white. However, looking at where this voter lives, we would be more likely to say that she is African American, with 96.39% of her Census Block being African American. After combining these two pieces of information, the updated probability of this voter being African American is improved to .9886.[6]

[6]Note that some probabilities in Table 1 are listed as 0. Zero probability events are, rightly, often viewed with suspicion. However, in this case, if the Census has perfect measurement and there is no movement, the probability is 0. Take, for example, the voter named Reid. There are no whites in her Census Block. If this count is correct, then the probability that she is white is certainly 0. Of more interest, however, is that the probability of her being Black is more certain than what could be determined by only looking at the Census Block demographics. Similarly, with the voter named Mogel, there are African Americans in her

With these estimates, I have a probabilistic determination of racial identity.[7] I use the notation $p([race]|name)$, where, in this case, *race* is *white* or *black*. This is shorthand for $p([race]|$name & location$)$ because the estimate is a function of both name and location.

3.2 2000 and 2004 elections

I measure the treatment effect by comparing turnout in the 2000 and 2004 General elections. Using the 2000 and 2004 General elections is desirable because this pair of elections involved a similar set of contests and came immediately before and immediately after most of the demolition of the projects. Also, Presidential elections are particularly appropriate for tests of group threat because of their national focus. Isolating the effects on behavior implied by group threat requires elections where voting behavior is unlikely to have been motivated by local issues or candidate contests that could turn in part upon the presence (or absence) of the housing projects or their residents. Notably, these elections did not involve Aldermanic races or any elections for any other citywide offices for which local factional politics would be more salient. There is little reason to believe that voters see the Presidential race as directly connected to the housing projects. Because the housing projects were not at stake in the Presidential election, it is unlikely that voters voted with the intention of influencing the future of the housing projects.

This means that the situation examined in this paper is similar to that studied by Key (1949). Key was measuring Presidential elections. It seems unlikely that the white voters were drawing a direct connection between their local population demographics and the outcome of the Presidential election. Furthermore, the Blacks in Key's study were virtually disenfranchised, so they did not present an actual *electoral* threat to the white voters. Instead,

Census Block, but her probability of being Black is 0 because the Census Bureau has counted no African Americans named Mogel in the United States. Even if the Census Bureau is wrong about this count of the name Mogel, the probability calculated here is probably a fairly accurate assessment of the likelihood that this particular individual is white or Black.

[7]Some voters live in Census Blocks that are perfectly racially homogenous, therefore the estimates generated by my updates have probability 1 that the voter is of a certain race. These estimates are almost surely wrong, due to measurement error by the Census and demographic change over time. However, even self-identification of race in surveys has associated error due to coding error and misreporting.

Key hypothesized that some Southern whites were habitually politicized by a psychological reaction to the presence of Blacks. Key describes this as "Negrophobia". Similarly, it seems that the white voters living near the Chicago housing projects were likely not motivated by influences directly related to the housing projects, but were instead politicized by the simple proximity of their Black neighbors.

4 Estimation and results

I want to measure the change in voter turnout for white voters that were treated by living in close proximity to the demolished public housing projects. This analysis can be thought of as an experiment in which the treatment is the demolition of the projects and in which the treatment group are white voters who live nearby, while the control group are white voters farther away. In most simple terms, the effect of the treatment is the difference in mean turnout at time t and $t-1$ for white voters close to the demolished projects relative to white voters not close to the demolished projects. I have to choose a distance from the housing project for which to measure the treatment effect. Call this distance d^*. If each voter lives d distance away from the nearest demolished project, then white voters for whom $d < d^*$ are the treatment group. Voters for whom $d > d^*$ are the control group. By taking the difference between t and $t-1$ for the control group, I account for the average change in turnout across the city. This is the change in turnout that was experienced, on average, by everyone in the city and is not attributable to racial threat. I subtract this control group difference from the treatment group's difference in turnout at t and $t-1$. This difference in difference yields the effect of the treatment (T). This is:

$$T = [P(Vote_{t-1}|d^* > d) - P(Vote_t|d^* > d)] - [P(Vote_{t-1}|d^* < d) - P(Vote_t|d^* < d)] \quad (1)$$

This is a very straightforward test: relative to the change in voting in the rest of the city, did white voters close to the project vote more or less after the projects were demolished? If

they voted less, then this might be attributable to racial threat. After reporting the results of this test, I will do several robustness checks to look at competing explanations.

There is no obvious choice of a cutoff distance, d^*. Some political phenomena have been shown to operate over a large spatial distance, for example across multiple counties (Gimpel, Karnes, McTague, & Pearson-Merkowtiz 2008), but it is not clear over what distance the treatment in this study should have an effect. Key's (1949) original insight about racial threat was within counties, which might be considered a large geographic space for a person to be affected by the demographics of one's neighbors. Key was also using only aggregate numbers, so it is not clear if the phenomenon he observed was due to certain white voters in the county – such as those in immediate proximity to African Americans – or whether all whites in the county were affected. According to spatial impact theory, the effect of the treatment should decline as the distance from the project increases. As an initial test, I will look at d^* at increasing distances from 100 meters to 1 kilometer from the projects.[8]

I also have to define voters' race based on my probabilistic estimates described above. For example, to identify a voter as white, do I only include voters where the estimated $p(white|name) = 1$ (that is usually voters that come from perfectly homogenous Census Blocks) or do I allow for voters that have a lower probability of being white? The lower $p(race|name)$, the larger the sample size, but the greater the chance that my estimates are contaminated by misidentification of the voter's race.

Because of this tradeoff, I will test my findings across a number of specifications of race, from only including voters where the probability is calculated as 1, down to .975. These probabilities, which are expected to contain no greater than 2.5% error, should provide a

[8]It is my opinion that, in terms of everyday human interactions, one kilometer is a considerable distance in an urban area. Looking at the map of Chicago, a city of relatively high density, moving 1 km can take a person from, for example, the Gold Coast to the South Loop - very different neighborhoods in sociological terms. And to get there, a person would cross through several other neighborhoods. Generally speaking, it seems that 1 km is beyond what a typical person would consider their "neighborhood". In fact, political science research often uses geographies like a Census Tract or Block Group to approximate a neighborhood. If these are good approximations of a neighborhood, than a 1 km area is usually far larger than a neighborhood. In Chicago, moving across 1 km would take a person through several Block Groups and, potentially, several Census Tracts. In this sense, it would not be surprising if the treatment strength varies significantly over 1 km.

large enough sample without having too many false-positives so as to cast doubt on the treatment effect.

4.1 Results

Figure 1 is a graphical display of the results of the difference in difference test from Equation 1. In this figure, $p(white|name) > .975$. The dots are the difference between turnout in 2004 and 2000 for voters near the projects minus the difference for other white voters far from the projects. Each dot is at a different distance, d^*, moving away from the projects from 100 meters to 1 kilometer. The dot itself is the mean, with the dotted lines that extend vertically representing the 95% confidence intervals of this estimate.[9]

The N of the treatment group is plotted next to each dot.[10]

Whites near the projects voted less than whites elsewhere in Chicago after the projects were destroyed. The effect diminishes as the distance of the voter from the demolished housing projects increases, with the effect approaching zero after 700 meters (about 4/10 of a mile or 3.5 city blocks in Chicago). These results are consistent with both hypotheses: the white voters appear to react to racial threat and the effect declines with distance, as is expected by spatial impact. The results are the same with higher levels of $p(white|name)$. Figure 2 is the results of the same test with $p(white|name)$ set to five different levels from .98 to 1.

When d^* is small and $p(white|name)$ is high, the N is small and the estimates are more uncertain. However, moving away from the projects, and when $p(white|name)$ is relaxed, so that N becomes larger, the estimates clearly achieve statistical significance.

[9]Some analysts may not choose to refer to these results as "estimates" because they represent the precise results of a single event. In that sense, there is no estimation involved: I know the precise effect of the treatment on the population used here. However, I am comfortable calling these results estimates because I am not using the complete population of voters, but rather a sample necessitated by data limitations. And, I also refer to the results as estimates because I believe I am measuring a general phenomenon in that of Racial Threat, which is exportable to other times and places with a likely similar effect.

[10]Because the control group at each d^* is all voters where $p(white|name) > .975$, and where $d^* < d$, the N is always very large. The smallest realization of the control group N in Figure 1 is 34,915, with a maximum of 157,494.

The effect is also substantively large: the average effect is generally greater than five percentage points at distances less than 500 meters.

4.2 Possible causes associated with living near the demolished projects, other than demographic change

Perhaps the effects that I am attributing to racial threat are not due to the demographic shock that accompanied the removal of the projects but instead were caused by some other unmeasured event particular to the neighborhoods where the demolished projects were located. If the effect is not caused by racial threat then we might expect to see the same change in the African American residents of these neighborhoods as we see for the white residents. Conversely, African Americans would presumably not be affected by the threat of other African Americans in the same way as whites, so, if racial threat is causing whites to vote less in 2004 than 2000, we should not see similarly reduced turnout among Blacks. If African American turnout decreased as well, it is difficult to attribute the decline in white turnout to the demographic change.

To test this, I measure how the event affected African Americans of the same neighborhood. And to make the test more robust, I look for African Americans who are similar to the white treatment group in terms of other covariates. To do this, I ran a nearest-neighbor matching algorithm to select a pool of African Americans that were from the same neighborhood as the white sample and were matched as closely as possible on other observable covariates (Ho, Imai, King, & Stuart 2007a, Ho, Imai, King, & Stuart 2007b).[11] The samples were balanced on the following: party, gender, distance from the nearest project, age, average number of votes cast in 1996 and 1998, and median income. All of those variables are available, at the individual level, from information provided on the voter file, except median

[11]Neighborhood is defined by the radius d^* around a given project. So, whites who lived, say, within 200 meters of the housing project Cabrini-Green were matched with African Americans who lived within 200 meters of Cabrini-Green (but did not live in Cabrini-Green itself, see Appendix). I use voters where $p(white|name) = .975$ or $p(black|name) = .975$ to identify the voters as white or Black and I test the difference in means for samples from 100 meters to 700 meters.

income which I measured at the Block Group level (see the Appendix).

After matching the treated white voters with a similar group of African American voters, I can compare the mean differences in turnout between 2000 and 2004. This allows me to compare two groups that are very similar in all aspects except for their race, which should condition their susceptibility to racial threat. The difference in difference for the white treatment group and the matched Black treatment group are displayed in Figure 3. The white treatment group is represented by the white dots and the matched Black group by the black dots. The results for the white treatment group are the same as in Figure 1. The N of each treatment group is the same because they are created from matched samples, but the Black control group is larger than the white control group. The difference between the Black difference in difference and the white difference in difference is displayed at the top of the figure. The results of this test support my hypothesis and may even strengthen it. For every choice of d^*, the effect on the white treatment group is far greater than the Black treatment group. In fact, in four cases, the turnout of the Black treatment group actually increased after the demolition of the projects.

One way of thinking about of the effect of racial threat is as the difference between the effect on the white and Black treatment groups. If the effect on Blacks does not include any influence by racial threat, then the turnout of Blacks might be the same as what would be expected from whites in the same neighborhood without any influence from racial threat. In this case, then racial threat is the difference between the estimate for Blacks and the estimate for whites. At some d^*, these effects are even greater than when the effect on whites alone is considered. For example, at 200 meters, the point estimate of the difference in difference in difference is 10 percentage points.

4.3 Confounding and race

Perhaps there are other factors confounded with proximity to a demolished project that affect behavior differently for white and Black voters. It might be that matching with

African American voters is a poor control for white voters because some unmeasured influence operates on African Americans and not on whites, or vice versa. For example, a political campaign in the neighborhood that particularly targeted African Americans and, therefore, stimulated their turnout. So, it is important to control for other influences while keeping race constant.

To check for this, I can, again, test the treatment group against a matched sample, but this time matching white voters against other white voters that lived near projects that were *not* demolished. Presumably, once they have been matched on available covariates, the difference between these two groups is due to the treatment of the demolition of the projects and the accompanying removal of outgroup threat.

After matching, I run linear regressions to control for the potential confounders and then examine the effect of treatment. Regression analysis also allows me to measure the effect of racial threat in the face of the covariates included in the model.[12] The coefficient estimates on treatment generated by regression on the matched samples are displayed in Table 2.[13] These estimates tell a story similar to the simple difference in difference test. The average effects are similar to those from the simple difference in difference. The estimates are the regression of a dummy variable for voting on a dummy variable for treatment, so the interpretation is simply in percentage points. For example, at 200 meters, the treatment reduced turnout by an estimated 13.55 percentage points. It appears that the racial threat effect is robust to the inclusion of the available covariates and, importantly, is not present for whites living near projects that were not demolished.

[12] As with matching with Black voters, the sample was white voters with $p(white|name) > .975$ and I tested it on distances from 100 to 700 meters.

[13] Complete results of the regression model are available by request.

4.4 Effects from habit disruption or other phenomenon associated with the project

The test on the matched data also helps to eliminate another competing explanation, that is that the destruction of the property itself, and not the accompanying demographic change caused the change in behavior. This may especially be likely, for example, if polling places were located in the projects. However, that the effect is different between African Americans and whites undermines this explanation. There is no obvious reason why a change in routine or other disruption caused by the destruction of the project would disproportionately affect whites. In fact, because of the likely reluctance of many whites to visit the public housing projects, it seems that if the polling location were located in the project itself, moving the poll elsewhere might actually increase white turnout.

There may also have been politically motivating factors that were connected to the projects themselves. For example issues of crime and property values. Again though, that the effect of demolishing the projects is different for African Americans and whites undermines this objection. There is little reason to expect that these issues would affect African Americans differently than whites, especially when the comparison is across balanced samples so that the comparison is between whites and African Americans that are demographically similar except for race.

4.5 Other mechanisms?

A potential challenge to the racial threat interpretation is that, prior to the demolition of the projects, the whites near the projects voted at close to the same rate as whites not near the projects. If racial threat were motivating whites near the projects, we might have expected them to vote at a greater rate than those not near the projects. For example, in 2000, for whites living within 500 meters of housing projects, turnout was about 83%, while turnout for whites in the rest of the city was also about 83%. If racial threat were

stimulating turnout, we might expect that whites living near the projects would be voting at higher rates prior to the demolition of the projects. If this were a true experiment and groups had been randomly assigned to treatment and control, than whites near the project should be voting more than whites farther away, prior to demolition. However, because the groups were assigned by "nature", there could be covariates that caused the treated white voters to vote at lower rates pre-treatment than they would have if racial threat was the only factor affecting their turnout. If racial threat were affecting turnout prior to the treatment, then after controlling for other factors, prior to demolition of the projects, there should be a relationship between proximity to the housing project and turnout.

I can test for this by regressing turnout in 2000 on the available covariates used above and on $log(distance)$ from the project. For white voters, there is a statistically significant and negative relationship between distance and turnout. In other words, moving further away from the projects, whites were less likely to vote. This is consistent with the predictions of spatial impact and racial threat. Also supporting the hypotheses is that using the same regression model for African American voters shows no relationship between distance and turnout (the point estimate is actually slightly positive, but small and not statistically significant). This indicates that proximity to the projects affected white turnout and not Black turnout, as would be expected if turnout was motivated by racial threat. As one additional test, I regressed turnout in 2004, rather than 2000, on the same covariates and distance from the demolished projects only. Here, if the projects were what was stimulating turnout and not some other factor, then the relationship with proximity should not exist after the projects are demolished. This is what happens. The regression on turnout in 2004 yields an estimated coefficient on log(distance) that is close to zero and not statistically significant.

4.6 Subject misidentification

Since the voters are identified probablistically, what if they are misidentified? Could misidentification be driving the results, especially as $p(white|name)$ decreases? This is unlikely, be-

cause misidentification would bias the results towards zero, not towards a larger effect. The same test on African Americans and Hispanics (results available upon request) shows that the average effect on these groups is closer to zero or positive than for whites in the treatment group. In the case of Hispanics, it is strongly positive. So, if non-white voters were misidentified as whites, this would be a conservative bias, making the estimates closer to zero. Also, even if the bias were away from zero, misidentification is very unlikely to account for the magnitude of the results seen here. Take the treatment group at $p(white|name) > .975$, for example. If 2.5% of the subjects are actually misidentified and everyone of those misidentified voted consistently with the hypothesis (that is 100% voted in 2000 and none voted in 2004), which is extremely unlikely, this still cannot account for the 5 percentage point effect seen here. The maximum that the treatment effect could be attenuated, again in the extremely unlikely case that all those misidentified voted in 2000 and none in 2004 would be $.05 - .025 = .025$. Of course, there is no reason to believe that misidentified voters would differ so dramatically in their behavior from correctly identified voters.

5 Discussion

I have presented evidence for group threat that that is relatively free of questions of endogenous confounding. The exogenous intervention in the racial context of the residents caused a clear change in their voting habits. Additionally, the strength of the effect decreased with distance from the project, which is consistent with spatial impact. The estimated effect of racial threat, approximately five percentage points, is substantively large in many elections. That the effect is substantively large, might be further evidence for spatial impact because a large effect would be expected when the stimulus is an outgroup that is large, close, and spatially concentrated, as were the residents of the public housing projects in relation to white voters.

5.1 Racial threat and other influences

This finding does not parse out how much of the effect of the residents of the housing projects on their white neighbors was due to race and how much was due to other dimensions of difference, like poverty. The outgroup in the housing projects were different from their white neighbors in many ways, of which race was just one. Could the effect have been different had the residents of the housing projects (implausibly) not been poor? Perhaps, although that the effect I measure here seems to operate exclusively on whites makes it more likely that race is the most important factor. If it were an outgroup defined by poverty, then we might expect to see the same effect on African Americans.

However, spatial impact is a general theory of groups, so that it could apply to an outgroup defined in any meaningful way. The strength of the effect of spatial impact may be smaller when other outgroups, besides racial groups, are involved, especially since race is such an important social categorization in the United States. However this effect should apply to groups that are large enough to be relevant, spatially concentrated, and close enough to matter. Certainly, the African American residents of the demolished Chicago housing projects fit these criteria. An extension of this finding would be to test the strength of the effect when an outgroup is spatially separated in a similar manner, but their difference is defined on a dimension other than race: perhaps class, religion, or sexual orientation.

5.2 Racial threat and contact

The normative implications of these racial threat findings can be troubling. It may be tempting to say that racial integration leads to hostility. This could be viewed as evidence against the long-standing and controversial Contact Theory (Allport 1954), which argued that contact between groups leads to reduced hostility. However, in considering the implications of these results, it is important to note that the populations in this study were probably not meaningfully integrated. The white voters in Chicago were threatened by a spatially proximate, yet concentrated and separate, out-group population. It is doubtful that these whites

and African Americans thought of themselves as a common neighborhood community. I doubt the residents of the affluent Gold Coast referred to themselves as "living near Cabrini-Green", one of the nearby public housing projects. In the case of some Southside projects, the white and Black populations were separated by a literal barrier in the Dan Ryan Expressway.[14] As such, we do not know how whites would have reacted to their African American neighbors, had they been meaningfully integrated. This speaks again to the importance of considering proximity and concentration, in addition to size, in the relationships of groups.

The experience of public housing in Chicago does not speak to what would have occurred if the populations were truly integrated. And, as I have argued, it seems that the spatial separation, yet close proximity of distinct populations may be the ideal conditions for threat attitudes to arise. Urban planners have largely reversed the planning philosophy that made for this distinct segregation by race and class, and, as a result, the new face of public housing in the United States are mixed income, public/private ventures. Future work in this area should continue to find points of leverage for studying attitudes and behaviors in these newly integrated populations.

5.3 How generalizable is this finding?

There is, of course, a concern that this is an isolated phenomenon, only applicable to residents of Chicago. Perhaps Chicago, a city of extreme segregation and occasionally marked racial tensions, is more susceptible to racial threat than other locations. I think it is likely that this phenomenon could be seen across many settings due to the persistence of white racial attitudes in the United States. This is a testable claim. An attractive feature of this design is that the remaking of public housing was a nationwide phenomenon. Future studies could look for variation across cities and, maybe, even individuals to see how individual characteristics, such as racial attitudes, are correlated with responses to demographic change.

Public housing has undergone transformation across the United States. Some public

[14] Urban historians have argued that freeway construction was a common tool of deliberate racial segregation (Keating 2001)

housing agencies were careful to document the experience of the residents who were relocated because of the renewal projects. Academics have recognized the unique nature of this opportunity for studying the displaced residents (for example, Gay (2009)). However, documenting the experience of those who were *not* displaced is more difficult because they were never enrolled in any program or subject to any intentional intervention - their neighborhood simply changed around them. As a result, there is little documentation of who these people even are. However, from the perspective of research design, the subject left behind might be the most interesting subjects associated with this phenomenon because of the exogenous nature of the treatment. Expanding the study of the effects of the destruction of housing projects to more cities would require the collection of, at least, voter lists and geographic locations of the housing projects in more cities.

If this data is available, studying this phenomenon has many attractive features, not only for replication of my findings, but for extensions as well. Among the attractive features is that this is a natural experiment on an almost unprecedented scale. Only once before has there been such a large-scale, systematic transformation of America's urban landscape - that was during the federally-driven urban renewal of the mid-Twentieth Century, especially the construction of intracity, federal freeways. The advantage to researchers now studying the renewal of public housing is that technology is available for the data collection and analysis necessary for this type of study. The basic tools of my study - electronic voter files and GIS did not exist in the 1960's.[15] The availability of this technology means that this finding is extendable to almost every major city in the United States. Doing so would demonstrate whether Chicago is anomalous or whether spatial impact is a more universal phenomenon.

The implications for this theory go beyond merely the experience of urban dwellers - although cities are likely where the effect is strongest and where the large majority of people in the United States live. The dimensions of spatial impact: size, proximity, and concentration

[15]However, even now, the record keeping on voter files as recent as years old is inconsistent, at best. To enable an extension of this or other studies an effort should be made to preserve voter files, even as new ones are made available.

are applicable to any location - as is evidenced by Key's (1949) finding in the rural South. However, spatial impact may be of particular importance to a large-scale phenomenon shaping the future of American urban life. As large, Latino immigrant populations continue to enter United States cities, the socio-political dynamic between the recent immigrants and the established population of African American city dwellers is not unlike the situation between whites and African Americans immigrants in the postwar era (see Enos (2010)). The interaction of African American and Latino groups will likely be an important element in shaping the future of American cities.

6 Appendix

6.1 Imputation of race from names

To impute the race of voters based on their names, I do the following (this discussion is taken from **?**. The probability that a voter on the voter list is a particular race or ethnicity can be determined by first looking at the probability that their surname is associated with a particular race or ethnicity according to U.S. Census counts of names by race.

The Census Bureau tabulates the frequencies of races/ethnicities associated with last names in the United States. For example, Smith, still the most common name in the United States, is 73.35% white, 22.22% Black, 0.4% Asian and Pacific Islander, 0.85% American Indian, 1.63% mixed race, and 1.56% Hispanic.

I take these frequencies for the name of each voter and combine them with the racial demographics of Census Blocks. I use a simple application of Bayes' rule to update the name's initial probability of belonging to a given racial group with the probability that a voter is of a certain race based on their Census Block demographics.

In applying Bayes' Rule, the probability that a person with a given surname, S, is of a particular racial group g_i is:

$$Pr(g_i|S) = \frac{Pr(S|g_i)p_i}{\sum_{j \in G} Pr(S|g_j)p_j} \tag{2}$$

where $Pr(g_i|S)$ is the probability that any given name belongs to racial group g_i, so that $Pr(S|g_i)$ is the probability that any given individual from racial group g_i has the name S. And p_i is the probability that any given individual living within a Census Block is of racial group g_i.

G is the set of all available racial groups. There are six racial/ethnic categories for which the Census collects surname counts. So, for example, to calculate the probability that a person with name S is African American, $j \in G$ would consist g_j = African American,

plus the five other available racial group designations from the Census that are not African American. These would be white, American Indian or Alaskan Native, Asian or Pacific Islander, two or more races, and Hispanic.

The Census only lists name frequencies for names for which occur at least 100 times in the Decennial Census. I discarded all voters for which I could not calculate a race probability based on their name. Because of this, I was probably more likely to discard voters that came from small ethnic or immigrant minorities. Some voters could also not have their name probability calculated because I could not match their address with a Census Block.

6.2 Sample Selection

Here I describe the technical details of how the sample was selected.

Using the 2004 Illinois voter file, I identified the race/ethnicity of each voter, as described in the previous section. However, to do this, I first had to match each voter with the demographics of their Census Block. I did this by writing a computer program that sent each unique address in the voter file to the Census Bureau website's address search service. This returned the Census Block for each address, which I then combined with the voter file. I then matched the Census Blocks with demographic data from Summary File 1 from the 2000 Census. I obtained the Block Group level income information in the same way. This process, in addition to discarding voters because $p([race]name)$ could not be calculated, led to a total of about 150,000 voters being discarded.

Each voter was geocoded by obtaining a latitude and longitude for each address from a geocoding computer program. I was not able to geocode less than 1% of addresses. However, each address represented, on average, about 2.5 voters, so about 2% of the original voters were discarded during this process. I obtained an archived GIS shape file of public housing in 2000 from the City of Chicago, Department of Innovation and Technology. This was valuable because the shape files provide the two-dimensional perimeter of the entire project, rather than just a one dimensional point like an address or geographic coordinate. Using

GIS software that is integrated with a Structured Query Language database, I measured the distance between every voter and the projects. This is where the two-dimensional feature of the GIS shape files is important because it allows me to identify voters for whom the distance from the project is zero - meaning they live or lived in the project itself. These voters were discarded from the sample. However, GIS data is sometimes inaccurate. GIS coordinates can deviate from their actual location on the Earth. In order to avoid including in the sample voters that lived in the projects but because of measurement error did not have a distance of zero from the projects, I created a boundary of ten meters around the perimeter of each housing project shape and discarded all voters within this boundary. The logic in this was that it would discard voters that actually lived in the projects but for which geocoding errors had placed them outside of the projects' GIS shape. This removed several hundred voters, identified as African American, from the sample. I think there is little danger of false positive error in this case. These projects were free standing structures, unattached to other buildings and, often, separated by streets. As such, there seems to be little reason why a person would reside within 10 meters of the project if they did not actually live in the project. Leaving these discarded voters in the sample makes little difference for the estimates of the treatment effect.

Finally, I discarded all voters that had moved or newly registered after January 1, 2000. This is to eliminate voters that might have moved into the area of treatment after the projects were demolished and to keep the same samples for 2000 and 2004. With all the discarded voters, the final sample was reduced from 1.2 million to about 850 thousand voters.

References

Allport, Gordon W. 1954. *The nature of prejudice*. Cambridge, MA: Addison-Wesley.

Authority, Chicago Housing. 2000. Chicago Housing Authority: Plan for Transformation: Improving Public Housing In Chicago and the Quality of Life. Technical report Chicago Housing Authority.

Blalock, Hubert M. 1967. *Toward a theory of minority-group relations*. New York: Wiley.

Blumer, Herbert. 1958. "Race prejudice as a sense of group position." *Pacific Sociological Review* 1 (1): 3–7.

Bobo, Lawrence, & Vincent Hutchings. 1996. "Perceptions of racial group competition: Extending Blumers theory of group position to a multiracial social context." *American Sociological Review* 61 (6): 951–72.

Brewer, Marilynn, & Norman Miller. 1988. *Eliminating racism*. New York: Plenum Press chapter Contact and cooperation: When do they work?

Carsey, Thomas M. 1995. "The contextual effects of race on white voter behavior: The 1989 New York City Mayoral Election." *Journal of Politics* 57 (Febuary): 221–28.

Cunningham, Mark K., Susan J. Popkin, Michael Eiseman, & Kadija Ferryman. 2005. "De Facto Shelters: Homeless Living in Vacant Public Housing Units.".

Edsall, Thomas Byrne, & Mary D. Edsall. 1992. *Chain Reaction: the impact of race, rights, and taxs on American politics*. Norton.

Ellison, Christopher, & Daniel Powers. 1994. "The contact hypothesis and racial attitudes among Black Americans." *Social Science Quarterly* 75 (2): 385–400.

Enos, Ryan D. 2010. "Racial Threat: Field-experimental evidence that a proximate racial outgroup activates political participation.".

Farley, Reynolds, Charlotte Steeh, Maria Krysan, Tara Jackson, & Keith Reeves. 1994. "Stereotypes and Segregation: Neighborhoods in the Detroit Area." *The American Journal of Sociology* 100 (3): 750–780.

Fitzpatrick, Kevin, & Sean Shong Hwang. 1992. "The effects of community structure on opportunities for interracial contact: Extending Blaus macrostructural theory." *Sociological Quarterly* 33 (1): 51–61.

Fossett, Mark A., & K. Jill Kiecolt. 1989. "The relative size of minority populations and white racial attitudes." *Social Science Quarterly* 70 (4): 820–35.

Gay, Claudine. 2004. "Putting Race in Context: Identifying the Environmental Determinants of Black Racial Attitudes." *American Political Science Review* 98 (4).

Gay, Claudine. 2009. "Moving Out, Moving Up: Civic Engagement and the Move to Opportunity.".

Giles, Micheal W., & Melanie A. Buckner. 1993. "David Duke and black threat: An old hypothesis revisited." *Journal of Politics* 55: 702–13.

Gimpel, James G., Kimberly A. Karnes, John McTague, & Shanna Pearson-Merkowtiz. 2008. "Distance-decay in the political geography of friends-and-neighbors voting." *Political Geography* 27: 231-252.

Glaser, James M. 1994. "Back to the Black Belt: Racial environment and white racial attitudes in the South." *The Journal of Politics* 56 (1): 21–41.

Ho, Daniel, Kosuke Imai, Gary King, & Elizabeth Stuart. 2007a. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15 (3): 199–236.

Ho, Daniel, Kosuke Imai, Gary King, & Elizabeth Stuart. 2007b. "Matchit: Nonparatmetric Preprocessing for Parametric Causal Inferece." *Journal of Statistical Software*.

Huckfeldt, Robert, & John Sprague. 1987. "Networks in Context: The Social Flow of Political Information." *The American Political Science Review* 81 (4): 1197–1216.

Keating, Larry. 2001. *Atlanta: Race, class, and urban expansion*. Temple University Press.

Key, V.O. 1949. *Southern politics in state and nation*. New York: Knopf.

Kotlowitz, Alex. 1992. *There Are No Children Here: The Story of Two Boys Growing Up in The Other America*. New York: Doubleday.

Latané, Bibb. 1981. "The psychology of social impact." *American Psychologist* 36 (4): 343–356.

Latané, Bibb, & Sharon Wolf. 1981. "The Social Impact of Majorities and Minorities." *Psychological Review* 88 (September): 438–453.

Leighley, Jan E., & Arnold Vedlitz. 1999. "Race, ethnicity, and political participation: Competing models and contrasting explanations." *The Journal of Politics* 61 (4): 1092–1114.

Leighley, Jan E., & Jonathan Nagler. 1992. "Individual and Systemic Influences on Turnout: Who Votes? 1984." *The Journal of Politics* 54 (3): 718–740.

Massey, Douglas S., & Nancy A. Denton. 1993. *American apartheid: Segregation and the making of the underclass*. Cambridge: Havard.

Matthews, Donald R., & James W. Prothro. 1963. "Social and economic factors and Negro voter registration in the South." *The American Political Science Review* 57 (1): 24–44.

Mutz, Diana C. 2002. "The Consequences of Cross-Cutting Networks for Political Participation." *American Journal of Political Science* 46 (4): 838–855.

Oliver, J. Eric, & Tali Mendelberg. 2000. "Reconsidering the Environmental Determinants of White Racial Attitudes." *American Journal of Political Science* 44 (3): 574-89.

Putnam, Robert D. 2007. "E Pluribus Unum: Diversity and Community in the Twenty-first Century The 2006 Johan Skytte Prize Lecture." *Scandinavian Political Studies* 30 (2): 137-174.

Quillian, Lincoln. 1995. "Prejudice as a response to perceived group threat: Population composition and anti-immigrant and racial prejudice in Europe." *American Sociological Review* 60 (4): 586–611.

Rieder, Jonathan. 1985. *Canarsie: The Jews and Italians of Brooklyn against liberalism.* Cambridge: Harvard.

Sigelman, Lee, & Susan Welch. 1993. "The contact hypothesis revisited: Black-White interaction and positive social attitudes." *Social Forces* 71 (3): 781–95.

Tajfel, Henri, M. G. Billig, R. P. Bundy, & Claude Flament. 1971. "Social categorization and intergroup behaviour." *European Journal of Social Psychology* 1 (2): 149–78.

Taylor, Marylee. 1998. "How white attitudes vary with the racial composition of local populations: Numbers count." *American Sociological Review* 63 (4): 512–35.

Voss, D. Stephen. 1996. "Beyond racial threat: Failure of an old hypothesis in the new South." *The Journal of Politics* 58 (4): 1156–70.

Welch, Susan, Lee Sigelman, Timothy Bledsoe, & Michael Combs. 2001. *Race and place: Race relations in an American city.* New York: Cambridge University Press.

Wolfe, Heath. 2006. The Office of Public and Indian Housing is Taking Action to Overse the SEction 202 Mandatory Conversion Program. Technical report Office of Inspector General, Department of Housing and Urban Development.

Wright, Gerald. 1977. "Contextual models of electoral behavior: The Southern Wallace vote." *American Political Science Review* 71 (2): 497–508.

7 Figures and Tables

Table 1: Sample probabilities: race given name and location

	Name		Census Block		Bayesian Probability	
name	%white	%black	%white	%black	white	black
Collins	73.92	21.70	1.71	96.39	0.0010	.9886
Mogel	95.38	0	72.71	9.43	.9883	0
Reid	63.24	32.01	0	94.59	0	.9816
Haile	51.48	41.10	64.74	2.10	.7624	.1121
Gorman	91.48	3.11	55.26	24.56	.8163	.0698

Five names drawn randomly from the voter file to demonstrate my method for identifying voters' race. The first two columns are the frequencies of the name by race group in the United States. The next two columns are the demographics of the Census Block in which the voter lives and the last two columns are the updated probability that a voter belongs to a each racial group.

Figure 1: Effects of treatment across distances

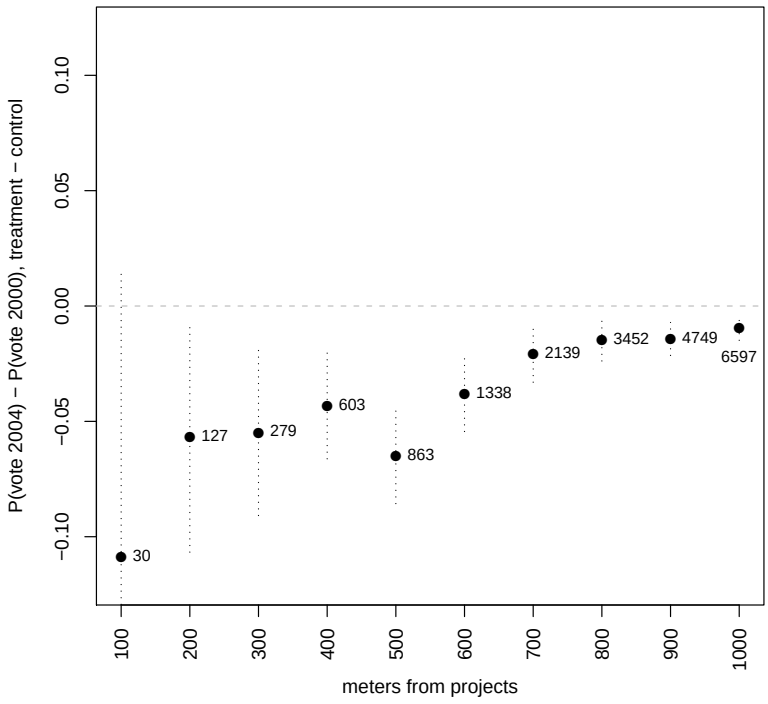

Difference in difference from 2000 to 2004 for treatment and control group. $p(white|name) > .975$. *The distance from the project on the x-axis goes from 100 meters to 1 kilometer. The treatment groups are all voters that live within each of these distances from the projects. Points below zero indicate that the treatment group voted less in 2004 than 2000, relative to the rest of the city. The N for the treatment group at each distance is printed next to the estimate. The dotted lines are the 95% confidence intervals of the estimates.*

Figure 2: Effects of treatment across distances with different $p(white|name)$

(a) $P(white|name) = 1$ (b) $P(white|name) > .995$ (c) $P(white|name) > .99$

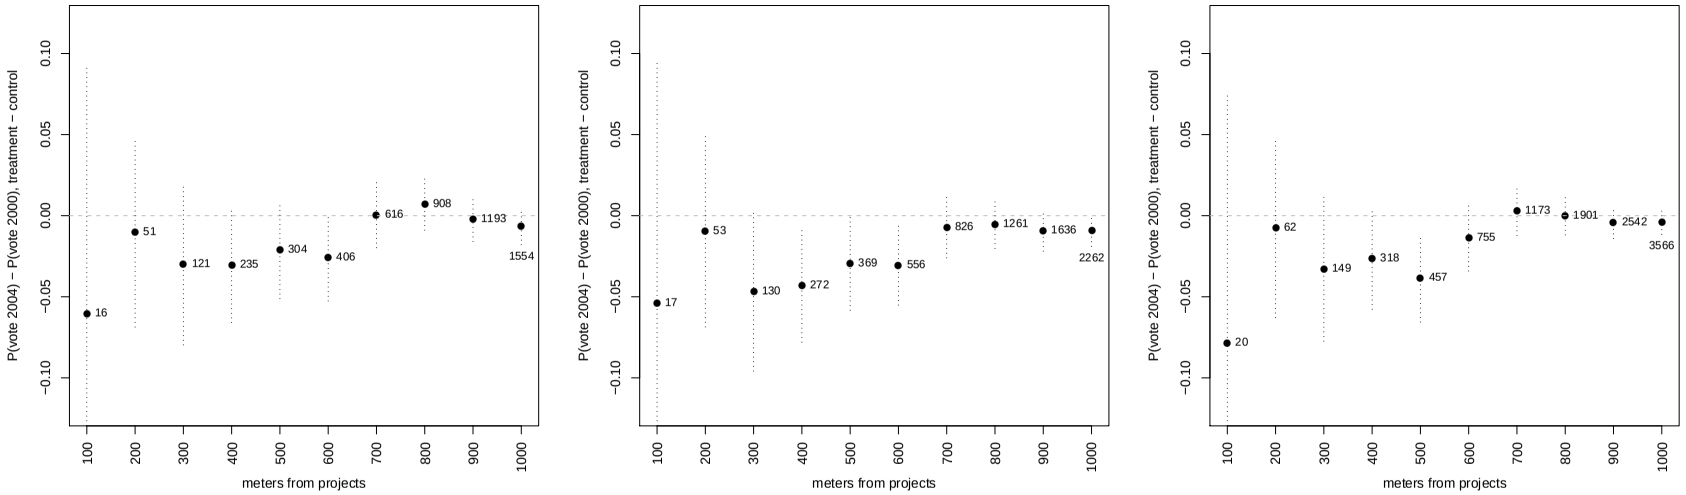

(d) $P(white|name) > .985$ (e) $P(white|name) > .98$

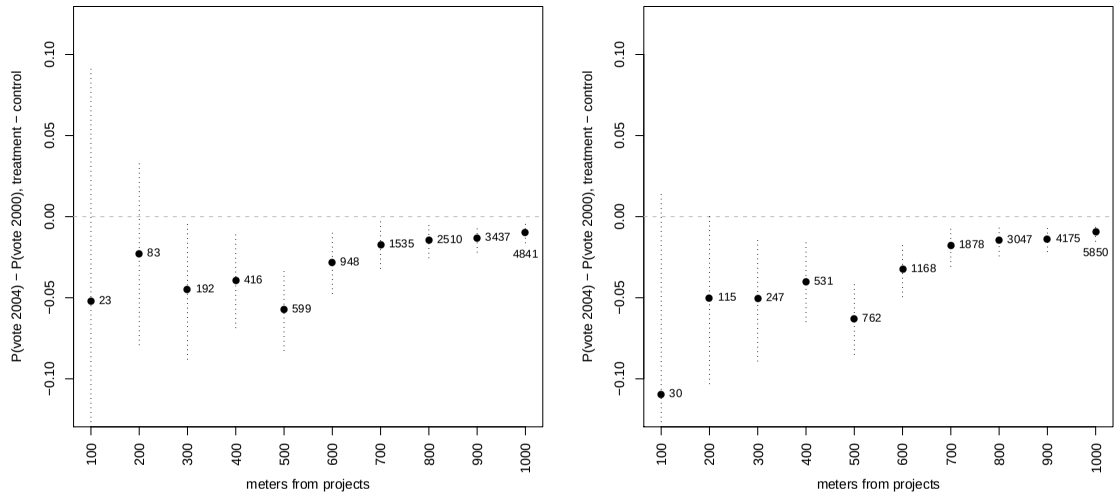

Difference in difference from 2000 to 2004 for treatment and control group. Each panel is for a different name threshold with distances from 100 meters to 1 kilometer. The treatment groups are all voters that live within each of these distances from the projects. Points below zero indicate that the treatment group voted less in 2004 than 2000, relative to the rest of the city. The N for the treatment group at each distance is printed next to the estimate. The dotted lines are the 95% confidence intervals of the estimates.

Figure 3: Difference between treatment on Blacks and whites

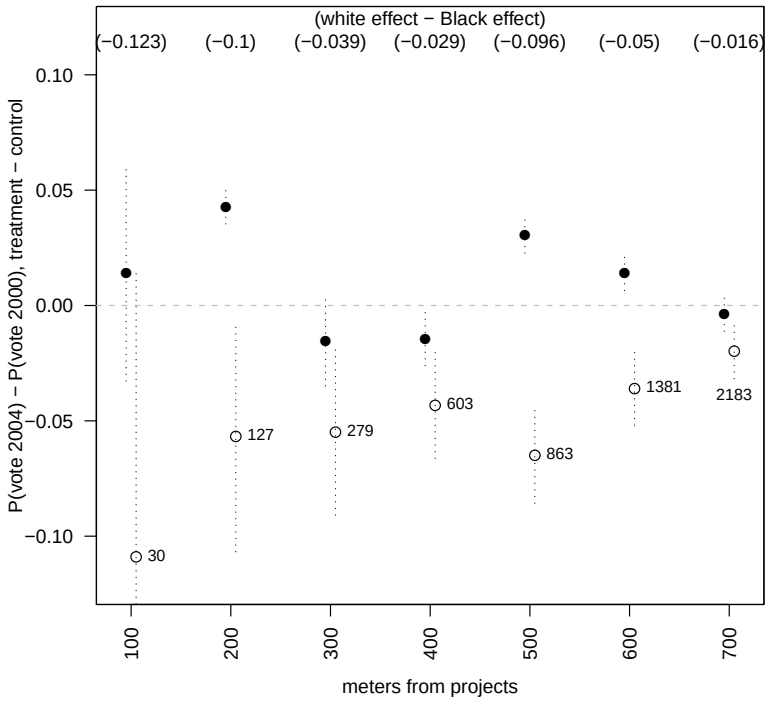

Difference in difference from 2000 to 2004 for treatment and control group for Black and white populations. White dots represent the white treatment group, black dots are matched sample of Black voters. Because they are matched samples, the N is the same for both treatment groups, although the N for the Black control group is larger in every case than the white control group. The N for the treatment groups is printed next to the white dots. The difference between the dots is plotted at the top of the figure. Points below zero indicate that the treatment group voted less in 2004 than 2000, relative to the rest of the city. $p(race|name) > .975$ *for both groups.*

Table 2: Regression estimates, effect of treatment on matched white sample

distance	β	SE
100	-.2478	.0916
200	-.1355	.0631
300	-.0406	.0464
400	-.0942	.0297
500	-.1211	.0241
600	-.0772	.0192
700	-.0498	.0150

Ordinary Least Squares regression estimates for the effect of treatment on matched samples of treatment and control group. Dependent variable is vote in 2004. Control variables are political party, gender, age, age^2, median income at the Block Group level, and turnout in 2000. Negative β means that the treatment group voted less than the control group in 2004.