How segregation drives voting behavior: New theory and evidence from South Africa*

Daniel de Kadt[†] Melissa Sands[‡]

This draft: December 29, 2015

Abstract

How does local demographic context shape political behavior? We investigate the effect of racial segregation on vote choice. We argue that racial segregation affects vote choice by shaping the racial context in which one lives, making racial identities salient in one's demographic context. We introduce an identity-centric theory of local effects, in line with situational theories of identity. We test the theory in the new South Africa by exploiting variation in the extent to which local segregation persisted after the end of Apartheid. Using a new panel dataset, which includes high resolution data from before the end of Apartheid, we estimate the effect of persistent segregation on voting in a instrumental variables framework. We develop an original instrument for segregation in the post-Apartheid period by using geographic measures of terrain hilliness. We then present individual level evidence, based on geo-referenced survey data for 42,000 people, that local racial context affects vote choice, as predicted by the identity-centric theory.

Keywords: Segregation, local context, voting behavior, South Africa

*Thank you to Matt Blackwell and Elliott Green for the detailed comments. Thank you to participants of MIT's African Politics Group, the Graduate Student Work in Progress seminar, Harvard's Political Economy Workshop, the Boston Working Group on African Political Economy, and the Harvard Meeting on Political Geography for valuable feedback. Thank you to Larry Zietsman and Izak van der Merwe for assistance with the 1991 census data, and to Helene Verhoef at Statistics South Africa for help with later census data. Thank you to James Aphane at the IEC for assistance with the 1999 election data. Thank you to the Human Sciences Research Council for making their survey data available, and to Benjamin Roberts for his assistance with that data. Thank you to Chris Wray and Samy Katumba of the Gauteng City-Region Observatory and Hayley Ivins Downes at Lightstone Property, for providing us access to transfer deeds data for Gauteng. The usual disclaimers apply.

[†]Department of Political Science, MIT, ddekadt@mit.edu.

[‡]Department of Government, Harvard University, sands@harvard.fas.edu.

1. Introduction

Segregation is a global demographic phenomenon. From the United States to sub-Saharan Africa, physical spaces partition groups along racial, religious, ethnic, or even economic lines. In some cases segregation is the deliberate result of state policy (Davies, 1967; Massey, 1993); in many more it is the result of endogenous "sorting processes" by which people, in choosing where to make their lives, select into ethnic or racial blocs (Schelling, 1971; Telles, 1992; Ihlanfeldt and Scafidi, 2002; Zorlu and Latten, 2009; Clark and Rivers, 2013). Recent research has investigated the origins of segregation (Van Kempen and şule Özüekren, 1998; Bayer et al., 2004; Bajari and Kahn, 2005), and the consequences of segregation, be they psychological (Allport, 1979), social (Oliver, 1999; Kasara, 2014; Bhavnani et al., 2014), economic (Massey et al., 1987; Boustan, 2013), or political (Ananat and Washington, 2009; Enos, 2014b). Other recent studies have explored, more directly, how racial or ethnic context moderates political behavior (Gibson and Claassen, 2010; Ichino and Nathan, 2013; Kasara, 2013).

In this study we investigate how racial segregation, which shapes the racial context in which voters live, affects vote choice. We introduce an identity-centric theory that links local demographic context directly to racial voting. Voting is the quintessential political act, and in African contexts is often tied to voters' identities (Posner, 2004; Lieberman and McClendon, 2013). In line with situational theories of identity, we argue that segregation affects voting behavior by affecting the salience of racial identities (Young, 1976; Bates, 1983; Horowitz, 1985; Eifert et al., 2010). We also draw on numerous works that study minority and immigrant communities, to propose that context effects are moderated by the relative status of race groups (Bonacich, 1973; Rumbaut, 1994; Marcuse, 1997; Sanders, 2002). We argue that for those in national minority race groups, living in relative racial isolation (social enclaves) *increases* the salience of their racial identity, in turn *increasing* the likelihood of voting for the party mostly closely associated with their race. By contrast, for those in national majority race groups, living in isolation *decreases* the salience of their racial identity, *decreasing* racial voting.

We test these predictions by assessing the effect of racial segregation on racial voting in post-Apartheid South Africa. We operationalize segregation as the degree of local racial isolation, adapting measures from Massey and Denton (1988). Similar to studies that leverage the built environment to study segregation (Ananat, 2011; Kasara, 2014), we use natural physical geography as an instrument for racial isolation – hills, valleys, and ridges that may act as buffers and break points that prevent racial mixing. Using two different measures of "hilliness" we instrument for levels of contemporary white isolation at the political ward level ($n \approx 2,900$). Consistent with our theory, we find that an increase in minority (white) isolation in South Africa strongly predicts voting against the majority-black African National Congress (ANC). We then present a ward-level difference-in-differences analysis to isolate the marginal effect of changes in the degree of racial isolation within political wards, while controlling for changes in the demographic levels of each ward. Again we find evidence that higher levels of racial isolation of the minority (whites) induce racial voting.

Using geo-referenced individual level survey data ($n \approx 42,000$), we then show that 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. Critically, we find the opposite for black South Africans. Black South Africans living in areas of high black isolation are *more* likely to vote across racial lines. In sum, these results accord with our identity-centric theory, but cannot be explained by available alternative theories.

While the issues that frame this paper are of global interest, the South African context is particularly suited to study. Apartheid, which lasted from 1948 through the early 1990s, is one of history's most infamous cases of state-sanctioned and -enforced racial segregation. Millions of black South Africans were physically displaced, dispossessed, and forced to live in cramped urban and rural ghettos. In Johannesburg, the most populous city in South Africa, black South Africans were forced to live 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 explicit acts of policy; the Apartheid regime planned its cities extensively, and used physical barriers like hills and ridges as pre-made buffers, insulating the white population from the larger non-white population (Davies, 1981; Christopher, 2001). While small numbers of black, coloured, and Indian South Africans were permitted to live in white-designated areas, racial mixing on any larger scale was rare and short-lived.

As the legislative framework of Apartheid was repealed in the run-up to the landmark 1994 elections, legal segregation came to an end. This precipitated a rapid return of black South Africans to previously white cities, surburbs, and peri-urban areas. Black South Africans, who had previously been forced to live in dense, under-resourced, and racially specific areas, crossed into the inner cities and suburbs in large numbers. This temporal shift around 1994 allows us to examine variation in the post-Apartheid period, conditional on Apartheid-era demographic profiles. We compare places with very similar Apartheid-era demographic profiles, and estimate the extent to which differences in levels of segregation in 2011, 20 years after the end of legal Apartheid, relate to voting behavior. To exploit this variation we develop a new dataset, a 1991 - 2011 panel that is the first to include high-resolution country-wide census data from the Apartheid era.

Our study extends previous work on the context effects of race and ethnicity. Prior studies have argued that sustained segregation may result in lower levels of inter-racial "contact", or, alternatively, an increased sense of "threat" between groups (Allport, 1979). More recently, social scientists have explored how exogenous changes in the racial composition of U.S. cities directly affects political behaviors (Ananat and Washington, 2009; Enos, 2014b), and similar arguments have been applied to parts of the developing world (Kasara, 2014). Research in African politics has proposed that segregation and racial/ethnic context shapes voting behavior through a rational interest mechanism (Ichino and Nathan, 2013; Kasara, 2013). Our study introduces a third complementary explanation: local context cues particular identities, which in turn cue particular behaviors. This approach integrates findings from the situational identity school of thought with psychological and rational approaches outlined above. We provide strong empirical evidence that segregation affects political behavior in a way consistent with this mechanism.

2. Theory

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 various 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.

Prior research has typically proposed two primary explanations that could account for how racial context may affect political behavior. The first explanation draws on Allport's (1979) early work on "contact" and "threat". Allport (1979) argued for two divergent effects of context. On the one hand, 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 these re-calibrated racial preferences then influence political behaviors, particularly when parties have racially-focused policies (for example affirmative action). This explanation is appealing, and Ananat and Washington (2009), Hopkins (2010), Enos (2014b), and Enos (2014a) find evidence in the U.S. that context may affect opinions and turnout. It is arguable, however, that theories of "contact" and "threat" are somewhat overfit to the U.S. case. That is, they appear to capture the effects of increased integration *for the white majority.* How to generalize beyond the U.S. is not obvious.

An alternative approach is introduced by Ichino and Nathan (2013), who make the link between context and *vote choice* more explicit. They argue that, in the Ghanaian case at least, local context affects voting behavior through rational self-interest. Assuming that voting is informed by expectations over targeted spending, ethnic (or racial) context informs beliefs 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. This argument makes a reasonable assumption: that individual voters consciously link race or ethnicity to distributive outcomes. 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. In these cases, it is unclear whether rational self-interest could explain context effects.

To fill these gaps we introduce a complementary explanation that draws on recent insights into ethnic voting in the study of African politics, and recent research that studies ethnic and migrant enclaves. Our theory explicitly links local racial or ethnic context to political behavior, and is applicable in non-competitive and highly centralized settings. In African political settings, as well as elsewhere, vote choice is often driven by ethnic and racial concerns. Yet situational theories of identity suggest that identities are multi-faceted, malleable, and subject to contextual activation (Young, 1976; Bates, 1983; Horowitz, 1985; Eifert et al., 2010). We argue that segregation, which shapes local context, may directly affect the *salience* of particular identities, influencing voting behavior. Crucially, we also consider how majority or minority status moderates context effects, proposing that members of minority and majority groups will act differently when living in isolated or mixed conditions, because racial identity is cued differently.

This identity-centric explanation is particularly attractive in cases where identity politics, be it racial or ethnic, is prevalent, and where voting along these lines is common. Our theory predicts that racial context, perhaps in addition to shaping preferences or beliefs (Ichino and Nathan, 2013), also serves to make salient racial identity, which translates into racial voting.

In settings with multiple race or ethnic groups of varying size, suppose that mixing occurs with some degree of variation – some places are mixed, while some are more homogeneous. We suppose that the effects of mixing vary depending on the degree of asymmetry in a society. For those in a national minority group, being surrounded by others from the same group places them in a monoracial sub-community where their minority status is highly salient. This proposition is supported by a large body of research on social identity theory (Tajfel, 1974), as well as scholarship that considers integration and alienation in immigrant communities (Bonacich, 1973; Rumbaut, 1994; Marcuse, 1997; Sanders, 2002). Social identity theory predicts that intergroup behavior is shaped by perceived membership in social groups. In racial or ethnic enclaves, the salience of group membership is magnified. For example, a panel study conducted by Sidanius et al. (2004), finds that membership in ethnic organizations augments minorities' sense of ethnic identity and perceptions of intergroup conflict.

Studies also suggest that that ethnic groups that live in physical enclaves, in which they are segregated or isolated from the rest of society, typically develop stronger ethnic identities, often through the reproduction of ethnic and linguistic identities in children. In many cases this does not represent an intentional redoubling of efforts of maintain their ethnic identites; rather, immigrants who live in enclaves tend to have lower levels of educational attainment, and are more likely to attain a high degree of fluency in their home language rather than the language of their host country (Borjas, 1995; Gang and Zimmermann, 2000; Chiswick and Miller, 2002, 2007). This, in turn, can further reinforce the group identities of members of the minority.

By contrast, for those in the national majority group, being surrounded by co-members does little to further racial identification on its own. Yet living in a mixed area, where interactions with members of the minority group are more frequent, may raise group salience, perhaps as a result of threat, inducing increased racial identification. Our theory simultaneously makes two strong predictions, over and above those made by either rational or threat based theories. First, minority group isolation should lead to *increased* racial voting. Second, majority group isolation should lead to *decreases* in racial voting. These dual hypotheses, illustrated in Figure 1, cannot be accounted for by currently available theories.

3. Data

Our theory is appealing in the South African case, a non-competitive democracy strongly divided along race lines, with highly asymmetrical race groups. The racial cleavage cuts socially, economically, and politically. Electorally, race acts as a powerful information heuristic for South African voters (Mattes, 1995; Ferree, 2006, 2010). The ANC represents a broad church for all black voters, and the political landscape offers few realistic alternatives to identity-based black voters. There are smaller ethnic or regional parties, but their quality, power, and appeal are limited. The ANC also carries a strong identity caché for black voters because of its status as the party of liberation. Many black South Africans describe the ANC as their "only

Figure 1: 2x2 of relationship between minority/majority status, isolation, and racial voting

Group Membership \ Racial Context	isolated	mixed
minority	racial voting increases	racial voting decreases
majority	racial voting decreases	racial voting increases

Note: For members of the majority racial group, isolation increases the salience of racial identity, thus increasing racial voting. Meanwhile, residing in a less isolated, more mixed area has the opposite effect. Members of the majority group experience the opposite: the salience of racial identity is activated in a more mixed, less isolated setting, thus increasing racial voting, and the salience of race decreases when the majority group is surrounded by members of her own racial group.

choice", and would rather exit the electorate than move to an alternative party.[1] This is particularly true in the period of our study, when the major alternative was the Democratic Alliance (DA), a white-led party that, while ostensibly liberal in nature, is typically seen to pander to the interests and demands of the white middle and upper classes.[2] These two parties soak up the vast majority of votes, and in 2009 and 2014, the DA and ANC together garnered over 80% of available votes. Similarly, in the most recent local election, the vast majority of local municipalities were won by either the ANC or the DA. Whites are thus associated with the DA, the party perceived to protect white interests.[3]

If our identity-centric theory applies, for white South Africans who live in areas in which whites are isolated, racial salience should *increase*, relative to those who live in more mixed contexts. Correspondingly, there should be an increase in DA voting in these areas, as white voters align themselves more readily with the white party. By contrast, black South Africans who live in areas in which blacks are isolated should find race less salient, relative to blacks who live in more mixed areas. Correspondingly, we expect a decrease in ANC support in these areas, as blacks vote less racially.

To test these claims we constructed a new and unique panel-dataset at the lowest political level in South Africa, the ward.[4] At the aggregate level, our interest is in estimating the relationship between racial segregation and voting behavior *for wards*. To do so we conduct a between-ward analysis. We measure hilliness, segregation, and political behavior at the ward level, and then make statistical comparisons across varying wards. The strategy is similar, though the unit of analysis more fine-grained, to Ananat's (2011) study of segregation in American 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 unit of analysis feasible.

For each ward, we combine three main sources of data. First, we source data on the physical geography of South Africa in a 1 km x 1 km grid, or raster layer, which is then aggregated up to the ward level. Second, we source census data from Statistics South Africa, 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 census, the final census conducted by the Apartheid regime, and gathered shortly before the end

[1]For instance, there was a major campaign during the 2014 national election called the "vote no" campaign, in which a number of older ANC stalwarts argued that voters should abstain from the election to protest against the ANC administration. Voting for the opposition was not an option, but exiting the electorate was (http://mg.co.za/article/2014-04-15-vote-no-say-anc-veterans-at-campaign-launch)

[2]There is now a new party called the Economic Freedom Fighters, but they do not figure in the period we study.

[3]All of our quantitative estimates are unchanged by using as the dependent variable "black party vote share" (ANC + Inkatha Freedom Party) or "white party vote share" (DA + Freedom Front).

[4]There are presently 4,277 wards in South Africa, electorally applicable for the 2011 local government and 2014 national elections.

of legislative Apartheid in mid-1991. We aggregated this data to the new 2011 ward boundaries. Third, we source electoral data for the National General Elections of 1999 and 2014, from the Independent Electoral Commission (IEC) of South Africa.

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, though we have data for 1996 and 2001, which we use in later analyses), within around 217 municipal clusters.

3.1. Measuring hilliness

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. The results are consistent across the measures. 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.[5] The *TRI* was developed by Riley et al. (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.

[5] 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.

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 $[\sum_{i=r-1}^{r+1} \sum_{i=c-1}^{c+1} (e_{i,j} - e_{r,c})^2]^{\frac{1}{2}}$. While TRI was originally developed to capture small-scale topographic heterogeneity, we apply them at a higher level of resolution. Having generated TRI 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. To do this we extracted two different summaries of the raster cells in each ward. First, we take the mean of TRI for each ward, and second, the variance of all values of TRI within each ward. We discuss the properties of these two measures below. All three of these ward-level aggregations (RangeAlt, MeanTRI, and VarTRI) are 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 markedly 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. Rather than measuring the hilliness of a given ward, it measures whether a given ward is *encircled* by hilly wards.

3.2. Demographic data

To measure the demographic properties of each ward we rely on census data collected and made available by the South African government. Our identification strategy relies on having data from before the end of Apartheid, which allows 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, nor has geographic data linking old censuses to physical space.

In 2014 Statistics South Africa (StatsSA) made available the 100% sample of various historical Apartheid censuses, conducted from 1960 to 1991. We acquired a 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 2011 ward boundaries. We then linked each observation in the officially released 100% sample of the 1991 census to the correct enumeration areas in 1991, and cross-checked our matches manually for accuracy. As a result we were able to match the 1991 individual-level census data with the 2011 wards, and aggregate the 1991 data from the individual level up to the ward level. This gives us a range of 1991 census measures at the 2011 ward level. The final dataset is novel in the study of South Africa, and includes population and socio-economic data at the 2011-ward level for 1991-2011. Variables and variable definitions are available in the appendix.

Limitations remain, however. The Apartheid government separated the country into independent Bantustans, four of which were considered "separate countries" by the Government. As a result, all censuses conducted before the end of Apartheid excluded the four independent Bantustans, mostly located in undesirable rural parts of the country. The 1991 data refers to roughly 2,900 of the 4,277 modern wards.[6]

Quantifying segregation in geographic space is a contentious issue in demography, sociology, and economics. We use *WhiteIsolation* (and, correspondingly, *BlackIsolation*) as our central measure of "segregation". We believe that these are good measures in the South African context because they reflect 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 in composition.

The *WhiteIsolation* variable was adapted from Massey and Denton (1988), and captures segregation *within* wards. To create this measure, we used fine-grained census data at the EA level for the years 1991 and 2011.[7] In any given 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

[6]There were separate censuses for three of the 10 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.

[7]The 1991 data is available at the individual level, clustered into EAs. For 2011 the data is actually available at the Small Area Level (SAL). This areal unit 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 the confidentiality of citizens. 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.

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 lives 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. We also calculate this measure for black South Africans, which we call *BlackIsolation*, and use in some individual-level specifications.

3.3. Election data

So that we can explore whether segregation has consequences for political behavior, we integrate into our dataset the election returns for the National General Election held in May 2014, again aggregated to the 2011 ward-level, which are the applicable wards for the period 2011-2015. We build these data into the ward-level panel outlined above. Unfortunately we do not have election data at the baseline, for the 1994 election, but do have data for the 1999 election.[8] 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. As a result the South African government has never made the disaggregated data available to researchers or the public, in an attempt to maintain the credibility of the first election in the eyes of the wider public. We use the election data we do have to generate voteshare variables, descriptions of which are available in the appendix.

3.4. Race and space in South Africa

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. The 1913 Natives Land Act cordoned off the majority of land in South Africa for whites, restricting property ownership and transfer rights for black South Africans. After Apartheid's formal inception in 1948, repeated legislation under the umbrella term "Group Areas" developed. This legislation divided the country's physical space into different areas for different groups. Each race or ethnic group was to reside in their allocated

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

areas – whites in wealthy, well resourced, large areas. Black South Africans, who made up roughly 70% of the population at the time, were to live in cramped, densely populated, and poorly resourced areas. The Group Areas Act of 1966 became the backbone legislation of this broader policy, aggressively regulating the right to live in certain areas, ownership and transfer of property, and even travel and commuting (Davies, 1981).

As the social and economic architecture of Apartheid began to collapse in the early 1990s, pressure mounted for a negotiated transition to inclusive democracy. In 1990 the National Party, under the leadership of F.W. de Klerk, released Nelson Mandela from Victor Verster prison, and lifted the ban on the African National Congress (ANC). In mid-1991 the government then undertook a series of legislative reforms, dismantling the legal architecture of Apartheid. In June 1991 the Group Areas Act was repealed, and the free movement of non-whites throughout South Africa was legally secured.[9] In 1994 the first inclusive democratic elections were held and the ANC, led by Nelson Mandela, was elected to power.

After the repeal of the Group Areas Act, black South Africans began to move out of the densely populated shanty towns and into different spaces within the cities and suburbs. As communities they faced collective choices about where to move and settle, constrained by the pre-existing arrangement of racial exclusion and ownership. In some cases they chose to settle in lower-income high-rise areas like Yeoville and Hillbrow in Johannesburg. In other cases informal settlements sprang up in the middle of cities; a famous example is the Cato Crest settlement in Durban, which lies between the white areas of Glenwood and Westville. By 1996, when the first post-Apartheid census was conducted, levels of segregation had declined greatly since the end of legal Apartheid in 1991, but remained marked and obvious nonetheless (Christopher, 2001).

The patterns of post-Apartheid settlement were in part determined by Apartheid zoning regulations from the 1950s. These regulations created "buffer strips" of empty space and used natural topographical features like ridges and hills to separate residential settlements of different race groups (Kuper et al. (1958), Davies (1981), Evans (1997), Christopher (2001)). By law, buffers prescribed to prevent contact between residential zones were a minimum of 200-500 yards, depending on the presence of roads and highways.[10] They were often much larger than this minimum, and were often created by displacing pre-existing communities and leaving land empty. These buffer zones thus remained empty and undeveloped during Apartheid, but were later settled as blacks rapidly moved into areas from which they were previously banned (Evans, 1997).

Figure 2 shows the distribution of race by ward from the 2011 census data for Johannesburg and sur-

[9] 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 mobility in the cities before 1991, but our data is the best we can access.

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

Figure 2: Johannesburg by race (2011) and hilliness

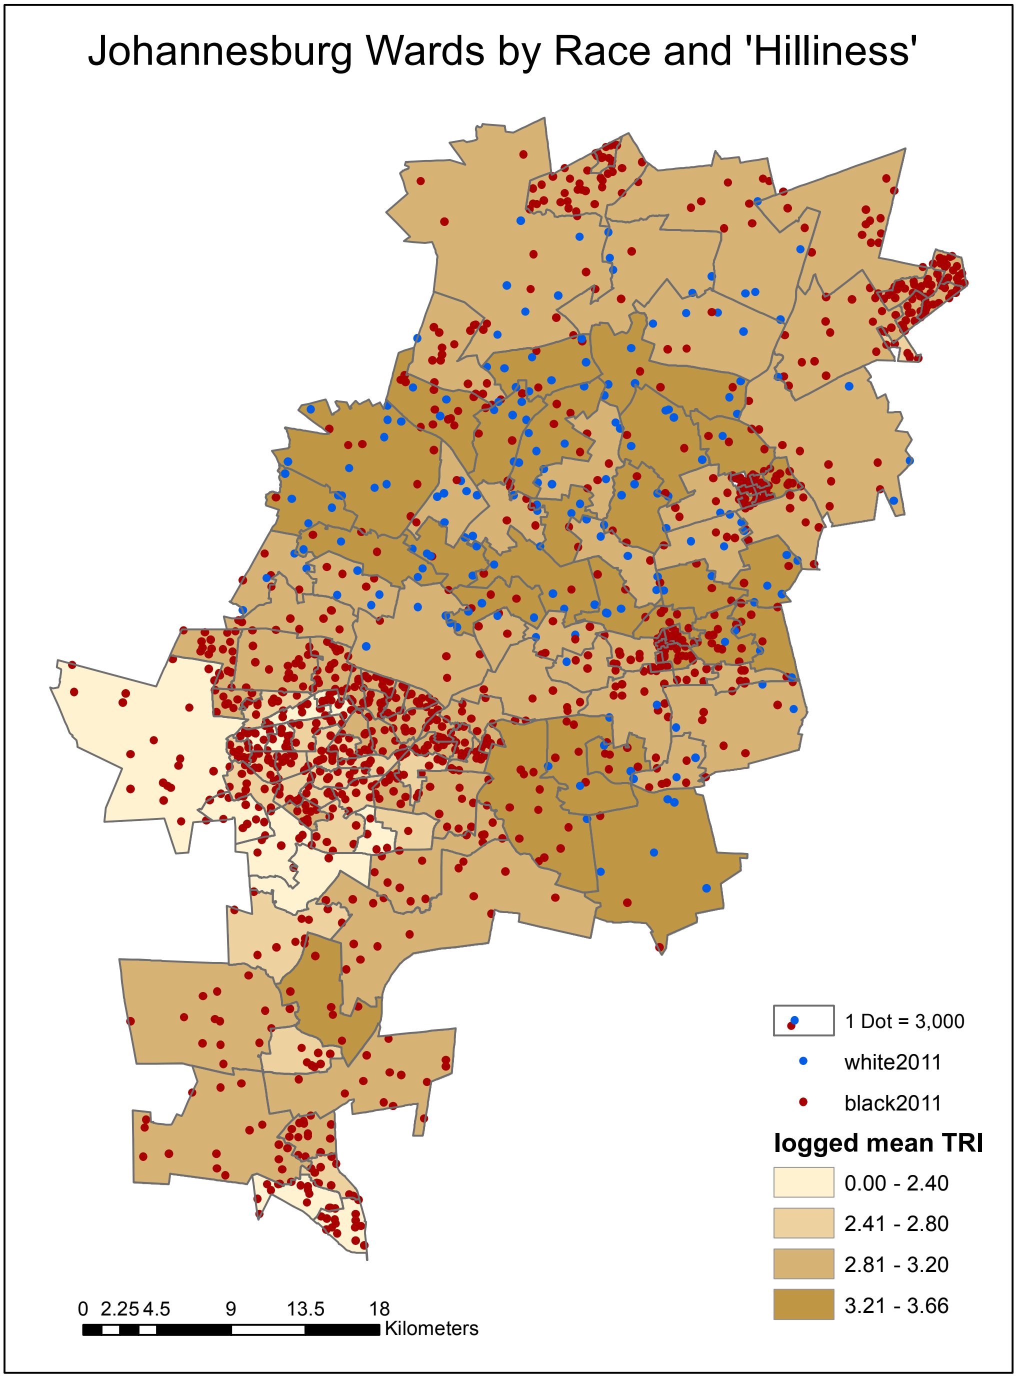

Note: This map depicts the city of Johannesburg. The chloropleth shading indicates values on *LogMeanTRI*, from low (light) to high (dark). The dots each represent 3,000 people. The blue dots are black citizens, and the red dots white citizens.

rounds. The chloropleth shading indicates values on one of our measures of hilliness – *LogMeanTRI* – from low (light) to high (dark). The dots each represent 3,000 people, with the blue dots representing black citizens, and the red dots white citizens. The map help to visually understand the relationship we hypothesize above. For Figure 2, the darker ridges running through the center of the city serve to segregate less densely populated, majority white areas, from more densely populated black areas. The dense mass of black population toward the bottom left is Soweto, while the mass in the middle right of the map is Alexandra. The dense mass slightly lower the middle is the old city center, and the associated areas of Hillbrow and Yeoville. Visually, it appears that re-integration patterns in post-Apartheid South Africa followed natural geography in Johannesburg.

We argue that these patterns of sustained segregation are due to post-Apartheid urban settlement patterns being constrained by the availability of undeveloped land. Natural barriers thus influenced post-Apartheid settlement in two ways. First, and most importantly, they partially determined the availability of space on which to build shanty towns and informal settlements. Thus, natural barriers continued to serve their Apartheid-era purpose: they circumscribed, to some degree, where black South Africans moving into the cities could settle. A second illustrative example is the settlement Cato Crest, an area in the Durban suburbs that was cleared in the 1960s by the Apartheid government to create a vacant buffer zone (Edwards, 1994). Given its political history, Cato Crest was identified as a prime location for resettlement and re-integration in the post-Apartheid period (Khan and Maharaj, 1998). In reality, the area soon became oversubscribed with informal housing, shack-dwelling, and limited social services. But the area could barely expand because it is squeezed between two large physical ridges, the main Durban ridge and the hills of Westville, both predominantly white areas. Thus white areas located in hillier parts of the city were insulated from demographic change. The fate was different for white areas lying in the flatter areas below the ridges. Areas like Manor Gardens and Glenridge, which lie before the ridges and abut Cato Crest, became rapidly more black than areas "protected" by natural geography, despite the fact that in 1991 they were demographically and economically comparable.

The Apartheid government expressly used natural geography as a tool to enable segregation, but with a very specific agenda (Davies, 1981; Christopher, 1987). The implicit unintended consequence, we argue, is that natural geography remains an important enabling feature in defining racial space. Of course, the extent to which physical space is partitioned by geography pre-dates any attempts to intentionally segregate populations. Yet the spatial distribution of race groups during Apartheid is no accident; the Apartheid government actively used physical geography to shape Apartheid cities. The way this distribution *changed*

Figure 3: Johannesburg by ANC vote share (2014) white isolation (2011)

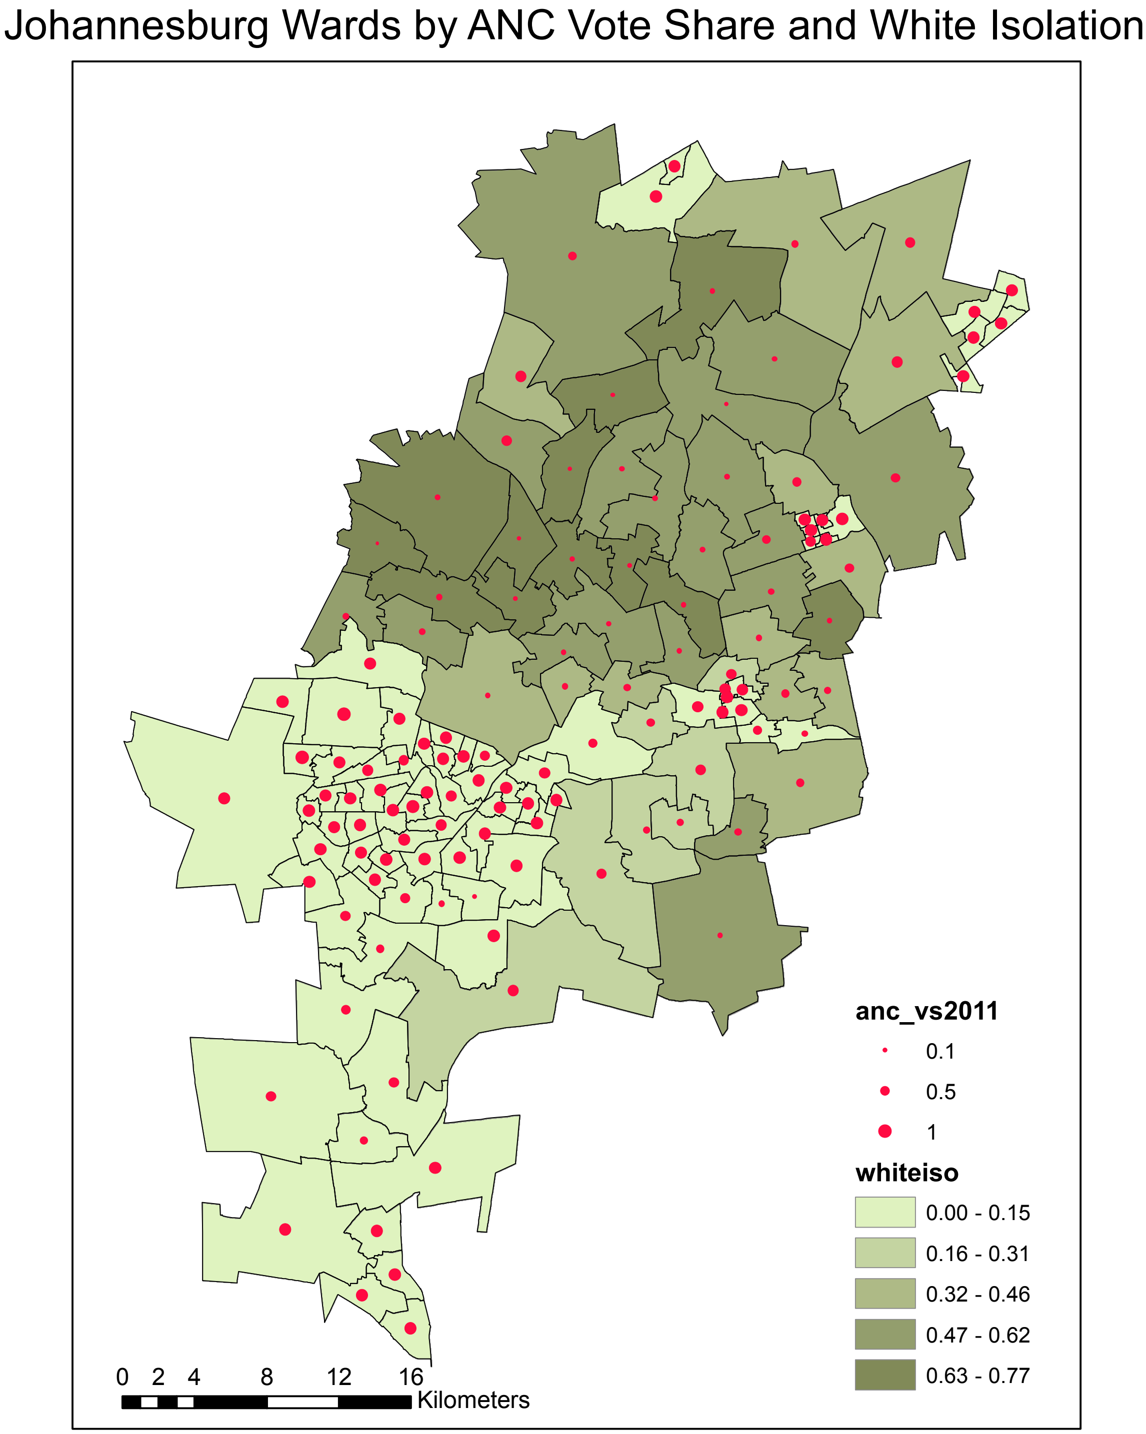

Note: This map depicts the city of Johannesburg. The chloropleth shading indicates values on *WhiteIsolation* in 2011, from low (light) to high (dark). The red dots located in the centroid of each ward are proportional in size to ANC vote share in the ward 2011. The largest dots appear in wards in which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

following the end of institutionalized segregation, and the degree to which that distribution persists today, is a function of both geography and baseline (Apartheid-era) segregation. Our research design exploits these facts.

Figure 3 shows exactly this correspondence. It maps the naive spatial association between segregation (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 centroid are sized proportional to the share of the 2011 vote that went to the ANC in that ward. Darkly 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. Contrasting Figure 3 with Figure 2 illustrates the logic of our research design. Areas with high hilliness have correspondingly high levels of white isolation, and lower levels of voting for the ANC.

4. Aggregate Level Analyses

We now explore this relationship statistically. Assessing the relationship between segregation and voting behavior is challenging, because 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. The IV approach is our primary research design, exploiting the context outlined in the previous section. We consider hilliness as a (continuous) "encouragement" – higher values of hilliness encourage higher levels of isolation or segregation in the post-Apartheid period. Segregation is a (continuous) "treatment", that has not been randomly assigned. We assess how this treatment affects political behavior, measured by ANC vote share. To do so we "instrument" for isolation with hilliness; this allows us to view 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. As a result, this approach allows us to isolate the effect of segregation, eliminating numerous potential confounders.

The IV approach hinges on two 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 Z affects our outcome of interest Y through no channel other than D. In this case it requires that hilliness not have an effect on political behavior outside of the effect channelled through segregation. In Section 6 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 regressions. The first stage specification is:

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

Where $Segregation_{i,2011}$ 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 $Segregation_{i,2011}$:

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

Where Y represents political outcomes of interest (vote choice). $\widehat{Segregation}_{i,2011}$ is the fitted value of $Segregation_{i,2011}$ 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 $Segregation$ 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.

4.1. IV results

The results of the first stage regressions can be found in Table 10 in the appendix. Both of our measures of hilliness induce meaningful variation in our measure of segregation, satisfying the relevance assumption of the instrumental variables identification strategy.

We present in Table 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, as shown in columns (5) through (8), to the inclusion of baseline economic covariates.[11]

Table 1: Instrumental Variables Regression of White Isolation (2011) on ANC Vote Share (2014)

	Dependent variable:			
	ANCvs2014			
	(1)	(2)	(3)	(4)
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)
Observations	2,866	2,898	2,850	2,882
R^2	0.690	0.704	0.688	0.710
Instrument	*LogMeanTRI*	*NN_LRA*	*LogMeanTRI*	*NN_LRA*
Municipal FE	✔	✔	✔	✔
Extra covariates			✔	✔

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

[11] It is possible that the result for black South Africans is driven by differential access to services. Once they move into majority white areas they receive a far higher level of government service provision than they may be accustomed to. As a result, they might vote ANC to reward them for this service. In regressions not presented here, we controlled additionally for the fraction of households in a ward that had formal dwelling structures, electricity for lighting, flushing toilets, and piped water, all measured in 2001. The results do not change.

To interpret the magnitude of these effects, consider the coefficient -0.690 in the first row of column (3) in Table 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.[12]

4.2. Difference-in-differences results

While the instrumental variables regressions provide evidence that higher levels of white isolation induce voting against the ANC, one might argue that the effects we estimate are merely evidence of racial voting. 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. While in the IV specifications we resisted including any covariates from after 1991, to avoid the possibility of post-treatment bias, we now present a difference-in-differences approach which 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, so we use $ANCvs1999$ as a proxy for the vote share in 1994. While an imperfect measure, it is the earliest available election data in South Africa.

The difference-in-differences specifications allow us to test whether more segregated areas vote less for the ANC, even conditional on changes in the fraction of the population that is white. It also has the added benefit of eliminating time-invariant or slow-moving confounders at the ward level. We approach this in two ways, seeking to explain change in ANC vote share at the ward level between 1999 and 2011. First, we simply take the change in white isolation from 2011 to 1991, 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

[12] 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%

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 a given change in the proportion of whites in an area affects ANC vote share varies by treatment status (levels of white isolation). Notwithstanding the limitations imposed by our proxy for 1994 ANC vote share, this approach eliminates time-invariant confounders, and identifies the effect of white isolation if there are parallel trends between treated and untreated units.

The results of this exercise are presented in Table 2, with the sixth and seventh rows of coefficients being the key difference-in-differences estimates. As predicted, 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). This implies that, 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.

5. Individual level evidence

The aggregate analyses are consistent with an identity-centric theory of context effects. At the ward level, it appears that, for a given sized white population, more segregated areas that are isolated from black South Africans, tend to "vote white" more readily. A key remaining test of the theory is whether *individuals*

Table 2: Difference-in-Differences Regression

	Dependent variable:			
	$\Delta ANCvs$			
	(1)	(2)	(3)	(4)
Δ *WhiteFrac*	-0.0644	-0.0488	0.00685	0.0128
	(0.0518)	(0.0517)	(0.0420)	(0.0437)
Δ *BlackFrac*	0.157***	0.169***	0.164***	0.172***
	(0.0399)	(0.0429)	(0.0399)	(0.0431)
Δ *ColFrac*	0.140***	0.151***	0.148***	0.156***
	(0.0403)	(0.0422)	(0.0398)	(0.0421)
Δ *WhiteIsolation*	-0.00260	-0.00205		
	(0.0191)	(0.0191)		
Δ *WhiteIsolationDummy*			0.000748	0.00226
			(0.00808)	(0.00825)
Δ *WhiteFrac*·Δ *WhiteIsolation*	-0.195**	-0.163**		
	(0.0755)	(0.0742)		
Δ *WhiteFrac*·Δ *WhiteIsolationDummy*			-0.146***	-0.118**
			(0.0506)	(0.0471)
Observations	2,824	2,809	2,824	2,809
R^2	0.552	0.552	0.552	0.552
Municipal FE	✔	✔	✔	✔
Extra covariates		✔		✔

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

residing in areas that are more or less segregated behave differently than their counterparts in more or less mixed areas. Survey data can provide a snapshot of political behavior in the various types of contexts. If segregation combined with the end of Apartheid generates persistence in ward-level voting patterns, we should expect to find evidence of those patterns among individual voters.

To examine context effects at the individual level, we spatially linked individual-level survey data from the South Africa Social Attitudes Survey (for 2004 - 2011) to the ward-level social, political, and economic variables described earlier. The resulting dataset contains 42,004 individuals over 7 periods.[13] We estimate the relationship between segregation on vote choice as a linear probability model and as a Logit, for whites and for blacks. The linear probability model (for whites) is estimated with the equation:

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

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 then repeat the exercise replacing $White_i$ with $Black_i$ in both the main effect and the interaction term. These specifications are estimated separately because of the collinear relationship between the $Black_i$ and $White_i$ dummy variables, but use the same underlying data. The results are presented in Table 3. Column (1) presents a regression without interaction terms, but with covariates, which offers a baseline estimate for white behavior.

In columns (1) through (3) of Table 3, the primary coefficient of interest is estimated on the interaction term $White_i * WhiteIsolation_w$. This tests the differential impact of racial context 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. Similarly, in columns (4) through (6), the interaction term $Black_i * BlackIsolation_w$ captures the marginal increase or

[13]In Table 3, we subset to voters only, about 60% of survey respondents. In addition, when we control for 1991 covariates the number of wards in the dataset drops considerably, leaving approximately 18,700 observations.

Table 3: Individual Interaction Results for ANC Voting

	Whites			Blacks		
	Dependent variable:			Dependent variable:		
	VoteANC			*VoteANC*		
	(1)	(2)	(3)	(4)	(5)	(6)
White	-0.802***	-0.772***	-0.735***			
	(0.0261)	(0.0264)	(0.0296)			
WhiteIsolation2011	0.0134	0.0251*	0.0407			
	(0.0137)	(0.0137)	(0.0259)			
White·WhiteIsolation2011	-0.110***	-0.126***	-0.149***			
	(0.0425)	(0.0424)	(0.0478)			
Black				0.887***	0.879***	0.896***
				(0.0226)	(0.0224)	(0.0252)
BlackIsolation2011				0.0809***	0.0786***	0.00267
				(0.0270)	(0.0261)	(0.0469)
Black·BlackIsolation2011				-0.0577*	-0.0753**	-0.121***
				(0.0323)	(0.0311)	(0.0380)
Observations	26,407	24,135	18,664	26,407	24,135	18,664
R-squared	0.478	0.491	0.497	0.478	0.491	0.498
Municipal FE	✔	✔	✔	✔	✔	✔
Indiv Covariates		✔	✔		✔	✔
Ward Covariates			✔			✔

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

decrease in black ANC votes due to the racial context in the ward.

Columns (1) through (3) suggest that whites living in greater isolation are especially likely to report voting against the ANC. Although whites in even 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 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.[14]

Figure 4: Marginal effects for ANC vote

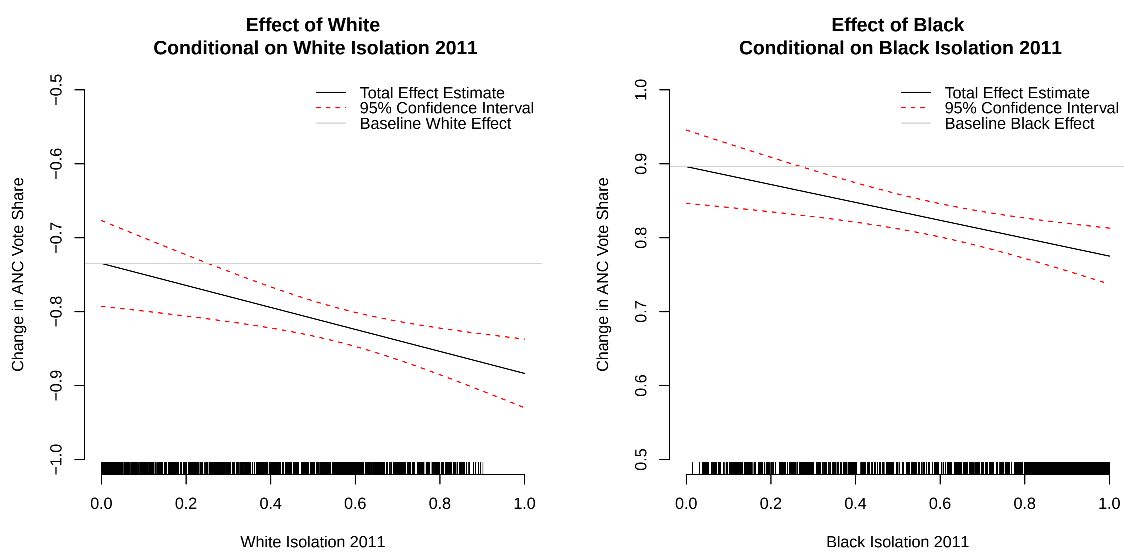

Note: These two panels plot the predicted vote for the ANC for each race group (white in the left panel, black in the right panel), as a function of ward-level racial isolation.

For blacks, columns (4) through (6) tell a different story. Blacks living in more mixed areas are actually *more* supportive of the ANC than blacks who live in areas of high black isolation. Here, we use $BlackIsolation$ as our measure of segregation because it describes how isolated a black individual is in her ward.[15] In more concrete terms, column (6) in Table 3 implies that, for black people, moving from the first quartile of $BlackIsolation$ to the third quartile of $BlackIsolation$ is associated with a 4.2 percentage point

[14]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.

[15]In general, blacks are more isolated than whites. The median ward has a black isolation index of 0.95, meaning that 95 of every 100 people a black person encounters will be black.

decrease in likelihood of voting ANC.[16]

This finding for black South Africans is critical evidence in favor of an identity-centric theory. While whites who are surrounded by other whites become less well disposed toward the black ruling party, blacks in less black areas actually consolidate their political preference for the ANC. This is consistent with our hypothesis that majority-minority status heavily moderates context effects. Black South Africans, as the majority group, move toward the ANC when surrounded by whites.[17] Whites, as the minority group, behave in the opposite fashion.

6. Discussion

Our analyses, at both the aggregate and the individual level, suggest that white South Africans vote "more white" when surrounded by fellow whites. Further, we find that black South Africans vote "less black" while surrounded by other blacks. These results are consistent across a range of research designs. The IV approach, for instance, eliminates numerous confounders by isolating "exogenous" variation in segregation. The difference-in-differences approach insulates our estimates from unit-specific time-invariant confounders. And the individual level results mean that our results are not driven by an ecological fallacy. Together, these results are consistent with an identity-centric theory, which takes into account how majority-minority status moderates context effects.

Crucially, we believe that these results cannot be explained by current theories of context effects. For instance, Ichino and Nathan's (2013) argument that context effects may be driven by rational self-interest cannot explain either our results for black South Africans (who become more likely to vote ANC in white areas) or the fact that context matters in a non-competitive setting. Crucially, this does not imply that prior research is incorrect, but that it is simply incomplete. Our theory is complementary: Context effects probably do involve self-interested rational voting, but may also involve identity voting.

Are there perhaps alternative explanations that may account for what we find? The primary challenge to our study is "sorting", which occurs when certain types of people choose to live in certain types of places. This would be a potential violation of the assumptions underlying our IV estimates. Our 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. We test for sorting in

[16]Blacks in the first quartile are 81.8 percentage points more likely to vote for the ANC, compared to their non-black counterparts, while black people in the third quartile are only 77.6 percentage points more likely to vote for the ANC.

[17]Keen followers of South African politics might suspect that our results are driven by KwaZulu-Natal, and in particular Zululand, where a large concentration of Zulus, who typically split their vote between ANC and IFP, live. To check that this was not driving our results we estimated the specification again without KwaZulu-Natal, and the results remain.

multiple ways, and ultimately assuage these concerns.

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, who are also less likely to support the ANC, chose to live in more hilly areas during Apartheid. Fortunately, 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, which are the focus of our analysis, it is less obvious that hilliness should be associated with wealth.

If particular types of individuals choose to live in more or less segregated areas, we anticipate associations between hilliness and economic variables like employment and education. Results are presented in appendix Table 5 that show very limited conditional associations between our measures of hilliness covariates in 1991. There are few statistically significant associations, and where they do exist the magnitudes are small. Where there is a conditional association the direction of the relationship actually suggests a conservative bias. Hillier places are slightly poorer (for one instrument) and slightly less educated (for both instruments). As a result, if hilliness is confounded, we should expect it to be positively associated with ANC vote share (as poorer and less educated people typically vote for the ANC), the opposite of what we find in our estimates. This suggests that the assumptions of the IV approach are likely met.

We also find stability in correlations between our instruments and 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 (economic) channels. If, however, the correlations remain stable over time, this is suggestive of limited change in observable covariates, and, hopefully, unobservables too. The results of this exercise are presented in appendix Table 6. 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.[18] 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.

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

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 socioeconomic 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 Table 7 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.[19]

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 (appendix Table 8). 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 (appendix Table 9). 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. In sum, our various robustness checks suggest that that sorting did not occur on a large scale, and cannot account for our results.

7. Conclusion

Prior research has considered the behavioral consequences of segregation, isolation, and mixing in both developed and developing settings. Extant theories have focused on two types of effects, rational or threat effects.

This study provides new insights into how one's lived context affects political behavior. In line with situational theories of identity, the identity-centric theory of local effects provides a mechanism for linking

[19]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.

racial or ethnic isolation and segregation to voting behavior. The identity-centric theory is complementary to existing work, but offers an holistic approach to racial or ethnic context effects for both majority and minority groups. The theory makes strong predictions; that more isolated members of a racial minority will find their racial identity more salient, while more isolated members of a racial majority will find their racial identity less salient. It further reconciles the situational school of identity, which has received great attention in the African politics literature, with theories of "threat" in the American literature. By proposing that context may shape self-identification, we make room for fluidity in perceptions of threat as a function of identity salience, a novel insight but one that is entirely consistent with new research on racial threat.

We tested the predictions of the theory at both the aggregate and the individual level in the context of South Africa. To do so we developed a new dataset of South African demography that spatially merges historical census data from 1991 to data in 2011. This ward-level panel data, in combination with an original instrument for segregation, allows us to estimate the effect of exogenous shifts in white isolation in the new South Africa. By comparing places that looked extremely similar on Apartheid-era demographic data, we establish the conditional ignorability of our instrument. This allows us to then leverage the exogenous shift created by the end of Apartheid in 1994, estimating how differential demographic effects of the end of Apartheid relate to contemporary voting behavior.

The results are consistent with the identity-centric theory – areas with greater levels of white isolation typically vote more racially. Even at the individual level, whites living in areas with higher white isolation engage in *more* racial voting, while blacks living in areas with higher black isolation vote *less* racially.

Apartheid was arguably the single largest act of state-led social engineering in human history, and our study provides new insights into its consequences. But the findings also have bearing outside of post-Apartheid South Africa. Scholars regularly point to essential similarities between Apartheid cities and other cities worldwide. For instance, much has been made of segregation in cities in the United States. We believe that this study has important implications for urban planners looking to restructure cities to encourage integration and diversity. Planners should be conscious of the fact that natural geography serves as an implicit set of barriers to integration, and that the effects thereof can be politically profound.

Our findings also show that these effects are heterogeneous by the status of race groups. From our perspective, this insight has important implications. Attempting to reduce racial voting by limiting the racial isolation of minority groups may have the corresponding effect of inducing higher levels of racialized behavior in the majority group.

References

Allport, G. W. (1979). *The nature of prejudice*. Basic books.

Ananat, E. O. (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, E. O. and E. Washington (2009). Segregation and black political efficacy. *Journal of Public Economics 93*(5), 807–822.

Bajari, P. and M. E. Kahn (2005). Estimating housing demand with an application to explaining racial segregation in cities. *Journal of business & economic statistics 23*(1).

Bates, R. H. (1983). Modernization, ethnic competition, and the rationality of politics in contemporary africa. *State versus ethnic claims: African policy dilemmas 152*, 171.

Bayer, P., R. McMillan, and K. S. Rueben (2004). What drives racial segregation? new evidence using census microdata. *Journal of Urban Economics 56*(3), 514–535.

Bhavnani, R., K. Donnay, D. Miodownik, M. Mor, and D. Helbing (2014). Group segregation and urban violence. *American Journal of Political Science 58*(1), 226–245.

Bonacich, E. (1973). A theory of middleman minorities. *American sociological review*, 583–594.

Borjas, G. J. (1995). Ethnicity, neighborhoods, and human-capital externalities. *American Economic Review 85*(3), 365–90.

Boustan, L. P. (2013). Racial residential segregation in american cities. Technical report, National Bureau of Economic Research.

Burchfield, M., H. G. Overman, D. Puga, and M. A. Turner (2006). Causes of sprawl: A portrait from space. *The Quarterly Journal of Economics*, 587–633.

Chiswick, B. R. and P. W. Miller (2002). Immigrant earnings: Language skills, linguistic concentrations and the business cycle. *Journal of Population Economics 15*(1), 31–57.

Chiswick, B. R. and P. W. Miller (2007). *The economics of language: International analyses*. Routledge.

Christopher, A. (2001). Urban segregation in post-apartheid south africa. *Urban studies 38*(3), 449–466.

Christopher, A. J. (1987). Apartheid planning in south africa: the case of port elizabeth. *Geographical Journal*, 195–204.

Christopher, A. J. (1990). Apartheid and urban segregation levels in south africa. *Urban Studies 27*(3), 421–440.

Christopher, A. J. (1994). *The atlas of apartheid*. Witwatersrand University Press Publications.

Clark, W. A. and N. Rivers (2013). Community choice in large cities: Selectivity and ethnic sorting across neighbourhoods. In *Understanding Neighbourhood Dynamics*, pp. 255–279. Springer.

Davies, R. (1967). The south african urban hierarchy. *South African Geographical Journal 49*(1), 9–20.

Davies, R. J. (1981). The spatial formation of the south african city. *GeoJournal 2*(2), 59–72.

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

Eifert, B., E. Miguel, and D. N. Posner (2010). Political competition and ethnic identification in africa. *American Journal of Political Science 54*(2), 494–510.

Enos, R. D. (2014a). Causal effect of intergroup contact on exclusionary attitudes. *Proceedings of the National Academy of Sciences 111*(10), 3699–3704.

Enos, R. D. (2014b). What the demolition of public housing teaches us about the impact of racial threat on political behavior. *American Journal of Political Science*.

Evans, I. T. (1997). *Bureaucracy and race: Native administration in South Africa*, Volume 53. Univ of California Press.

Ferree, K. E. (2006). Explaining south africa's racial census. *Journal of Politics 68*(4), 803–815.

Ferree, K. E. (2010). *Framing the race in south africa: the political origins of racial census elections*. Cambridge University Press.

Gang, I. N. and K. F. Zimmermann (2000). Is child like parent? educational attainment and ethnic origin. *Journal of Human Resources*, 550–569.

Gibson, J. L. and C. Claassen (2010). Racial reconciliation in south africa: Interracial contact and changes over time. *Journal of Social Issues 66*(2), 255–272.

Hopkins, D. J. (2010). Politicized places: Explaining where and when immigrants provoke local opposition. *American Political Science Review 104*(01), 40–60.

Horowitz, D. L. (1985). *Ethnic groups in conflict*. Univ of California Press.

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

Ihlanfeldt, K. R. and B. 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.

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

Kasara, K. (2014). Does local ethnic segregation lead to violence?: Evidence from kenya.

Khan, S. and B. Maharaj (1998). Restructuring the apartheid city. In *Urban Forum*, Volume 9, pp. 197–223. Springer.

Kuper, L., H. L. Watts, and R. Davies (1958). *Durban: a study in racial ecology*. J. Cape.

Lieberman, E. S. and G. H. McClendon (2013). The ethnicity–policy preference link in sub-saharan africa. *Comparative Political Studies 46*(5), 574–602.

Marcuse, P. (1997). The enclave, the citadel, and the ghetto what has changed in the post-fordist us city. *Urban affairs review 33*(2), 228–264.

Massey, D. S. (1993). *American apartheid: Segregation and the making of the underclass*. Harvard University Press.

Massey, D. S., G. A. Condran, and N. A. Denton (1987). The effect of residential segregation on black social and economic well-being. *Social Forces 66*(1), 29–56.

Massey, D. S. and N. A. Denton (1988). The dimensions of residential segregation. *Social forces 67*(2), 281–315.

Mattes, R. B. (1995). *The election book: Judgement and choice in South Africa's 1994 election*. IDASA, Public Information Centre.

Nunn, N. and D. Puga (2012). Ruggedness: The blessing of bad geography in africa. *Review of Economics and Statistics 94*(1), 20–36.

Oliver, J. E. (1999). The effects of metropolitan economic segregation on local civic participation. *American Journal of Political Science*, 186–212.

Posner, D. 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.

Riley, S. J., S. D. DeGloria, and R. Elliot (1999). A terrain ruggedness index that quantifies topographic heterogeneity. *intermountain Journal of sciences 5*(1-4), 23–27.

Rumbaut, R. G. (1994). The crucible within: Ethnic identity, self-esteem, and segmented assimilation among children of immigrants. *International Migration Review*, 748–794.

Sanders, J. M. (2002). Ethnic boundaries and identity in plural societies. *Annual Review of Sociology*, 327–357.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology 1*(2), 143–186.

Sidanius, J., C. Van Laar, S. Levin, and S. Sinclair (2004). Ethnic enclaves and the dynamics of social identity on the college campus: the good, the bad, and the ugly. *Journal of personality and social psychology 87*(1), 96.

Sparks, A. H. (1990). *The Mind of South Africa*, Volume 155. Heinemann London.

Tajfel, H. (1974). Social identity and intergroup behaviour. *Social Science Information/sur les sciences sociales*.

Telles, E. E. (1992). Residential segregation by skin color in brazil. *American Sociological Review*, 186–197.

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

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

Young, C. (1976). *The politics of cultural pluralism.* Univ of Wisconsin Press.

Zorlu, A. and J. Latten (2009). Ethnic sorting in the netherlands. *Urban Studies 46*(9), 1899–1923.

Appendix

Table 4: 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.
BlackIsolation	The proportion of black people who reside in the neighborhood of an average black person. Higher values indicated higher levels of black-centric segregation.
BlackFrac	The count of black individuals divided by the total population
WhiteFrac	The count of white individuals divided by the total population
ColFrac	The count of coloured individuals divided by the total population
IndianFrac	The count of Indian individuals divided by the total population
Income	The average income, from a 17-point discrete scale
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 5: 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 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 7: 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 8: 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 9: 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 10: 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