

Segregation, Spillovers, and the Locus of Racial Change

Donald R. Davis, Matthew Easton, and Stephan Thies*

July 11, 2024

Abstract

What explains the principal patterns and changes in racial segregation in American cities? Where does change occur and how shall this shape the way that we think about neighborhood racial tipping? Here we develop a discrete choice approach to neighborhood selection and racial segregation, but crucially relax the common assumption of zero spillovers of racial tastes across locations. Our approach allows us to nest Schelling's (1971) bounded neighborhood and spatial proximity models in general equilibrium. In simulations, we investigate the equilibrium properties of our model and derive hypotheses that we explore empirically using US census data from 1970 to 2000. We show that patterns in the data support the presence of spatial racial spillovers and show that these are key to understanding where neighborhood tipping, which we define as drastic racial change, is taking place. Our results on the locus of racial change conflict strongly with those of Card et al. (2008). For this reason, we revisit their work. We show that a combination of modern regression discontinuity methods plus theory-consistent restrictions lead their tipping results, which they define as dramatic white exit at critical thresholds of the minority share, to become precisely-measured zeros. Instead of tipping driving their results, we provide evidence that biased White suburbanization, closer to the ideas of Boustan (2010), explains the patterns they identify. In sum, our approach provides a distinct way of interpreting neighborhood tipping, as well as a promising path for understanding both the cross section and dynamics of spatial racial patterns in the data.

*Davis: Columbia University and NBER (email: drdavis@columbia.edu); Easton: Columbia University (email: me2713@columbia.edu); Thies: Columbia University (email: stephan.thies@columbia.edu). We are grateful to Jonathan Dingel, Tomasz Michalski, Dan O'Flaherty, and Allison Shertzer, as well as seminar participants at Columbia University and the London School of Economics. We thank Tommy Huang for his research assistance. We especially thank Iain Bamford and Howard Zhang for their contributions to earlier versions of the project. Easton acknowledges the Program for Economic Research at Columbia and the Alliance Doctoral Mobility Grant from Columbia and Sciences Po for support.

1 Introduction

In recent years, there has been a flourishing empirical literature on how neighborhoods, and in particular their socioeconomic and racial composition, affect life opportunities.¹ This includes randomized studies through the moving to opportunity experiments that provide credible causal evidence of these links (Chetty et al., 2016), and long-run observational studies that identify large causal effects of neighborhood residence on life outcomes in a quasi-experimental setting (Chetty and Hendren, 2018a,b; Chyn, 2018; Chyn et al., 2022). Naturally, the new empirical evidence on the importance of neighborhoods invites consideration of a variety of policies that might directly target social mixing or have indirect implications for the social composition of a location. Once we move past observational studies and measure-zero experiments to policy interventions at scale, it becomes crucial to think of general equilibrium reactions to any such interventions.² In a world in which drastic racial change may arise in response to local housing supply shocks, construction of infrastructure, and other policy interventions, there is a risk that endogenous sorting responses may partly or wholly undo the beneficial outcomes intended.

Racial residential segregation is among the most studied topics in the social sciences, and so one might imagine that the dynamics of segregation are well understood empirically and this understanding could inform policy. This, however, is not the case. Consider the foundational theoretical work of Schelling (1971), which has roughly 7,000 Google Scholar citations. Schelling developed there the bounded neighborhood and spatial proximity models. The first considers the fate of a single neighborhood and is the setting for articulating the highly influential tipping model. The second, often referred to as the checkerboard model, considers in an abstract setting an entire metropolitan community. This model illustrates a crucial spatial aspect of racial preferences, namely that even modest micro-incentives spilling across space can have large implications for macro-features and lead to high levels of segregation. Despite the broad academic success of Schelling’s two models, they have received relatively little formal empirical scrutiny in economics. The tipping model is examined in the influential work of Card et al. (2008), and ideas of tipping also play a role in a more recent structural literature on neighborhood and school sorting (Caetano and Maheshri, 2017; Blair, 2023). Racial spatial spillovers, however, receive little attention in formal empirical research, despite their potential relevance for understanding patterns of residential segregation and their implications for urban policy.

¹Chyn and Katz (2021) review this literature.

²An example of the type of general equilibrium evaluation we have in mind is Davis et al. (2021).

Contributions

This paper makes a number of contributions. The first is to develop a static model of discrete neighborhood choice featuring spatial racial spillovers. In doing so, we show how to nest Schelling's bounded neighborhood and spatial proximity models in a common, general equilibrium setting. Our integrated approach provides a simple, yet rich framework within which to think about neighborhood tipping and the formation of racial clusters in cities. We investigate the key equilibrium properties of our model in a series of simulations. Based on the simulation results, we articulate hypotheses that we then go to test using US census data from 1970 until 2000. Our hypotheses concern the racial composition of census tracts, their racial clustering across space, associated implications for rental prices, as well as the locus of racial change and the likelihood of tipping in response to local policies. In our analysis we focus on a dichotomous classification of race comparing non-Hispanic White households and minority households, understood to be the complement. For simplicity and following most of the existing literature dealing with the same time period from 1970 to 2000, we refer to the two groups as White and Minority households respectively.

A key feature of our model is the presence of multiple equilibria when racial preferences are sufficiently strong and extend beyond a single tract. To summarize predictions of our model in a way that is robust to hysteresis and to investigate our hypotheses empirically, we reduce the dimensionality of the US census tract data in a new way. We first operationalize the idea of clustering by considering sets of N (varying) or more contiguous census tracts with the same racial mode as a cluster. This permits us to investigate the salience of clusters in the racial geography of US metro areas. We then develop a novel way of visualizing the data. We focus on the distance in units of tracts of a location from a cluster boundary, i.e., a border where two tracts of different modal race abut. These visualizations allow us to examine key contrasts in the interior versus the edge of racial clusters, both in the cross section and across decades.

Our empirical analysis confirms that racial clusters are a ubiquitous feature of US cities over the period from 1970 to 2000. We show that there is a strong, but not precipitous gradient in racial composition at the boundary of racial clusters that slightly smooths out over time. We find that across all censuses, there is a persistently high degree of segregation interior to Minority clusters, even as there is a secular decline in the initially strong degree of segregation within White clusters. Likely reflecting differing degrees of homophily as well as wealth and income differences, prices are approximately flat within Minority clusters, but rise strongly as we move toward the interior of White clusters. Looking across periods of time, we find that drastic racial change occurs at the boundary of clusters, and the more

stringent our definition of drastic change, the more strongly that change is concentrated at the boundary.

Our last finding on the locus of racial change directly conflicts with prominent results reported by Card et al. (2008). In their reduced form investigation of neighborhood tipping points in the US, which focuses on the same time period as our analysis, the authors argue strongly against what they term the “expanding ghetto model”. Indeed they claim that tipping is strongest in areas remote from the existing Minority clusters. To see why we end up with such conflicting claims, we re-examine their results.

We first investigate the case of Chicago, which stood out as an extremely powerful example of tipping in the analysis of Card et al. (2008). They identify a tipping point between 1970 to 1980 at a tract’s Minority share of 5.7%. Moving beyond this tipping point leads to drastic White *exit* from a tract according to Card et al., with a discontinuous drop of roughly 30 percentage points in White population growth. A closer look at their data reveals a different story. Motivated by our more general equilibrium view of neighborhood change we find that their results from Chicago in 1970 instead seem to be driven by White *entry* to initially low Minority share, low density tracts in suburban areas. Rather than evidence for a tipping hypothesis, this is instead evidence for what we term “biased White suburbanization.”

We also revisit the regression discontinuity models employed by Card et al. to establish the significance of tipping points in different decades and across all MSAs. Our re-examination has two points in mind. The first is that preferred methods for regression discontinuity analysis have evolved, leading us to implement local linear, as opposed to global, polynomial regressions. The second is that since the tipping model is about White exit, not entry, a cap on observations of tracts experiencing spectacular total population growth should not eliminate the estimated tipping jumps. We implement this first by using population weighted regressions and then by excluding tracts with uncommonly high total population growth rates. Modern methods on their own eliminate the significance of the tipping results in 1980-1990 and 1990-2000, although these effects are imprecisely estimated. When we employ modern methods and add theory-consistent caps on total tract population growth to not be greater than 60%, all of the originally significant estimated tipping jumps from Card et al. (2008) become precisely measured zeros. The tipping results are gone.

Despite our finding that Card et al.’s (2008) tipping results were driven by biased White suburbanization, we argue that tipping defined instead as drastic racial change is an important real-world phenomenon. Instead of being an aspatial process, we highlight that spatial spillovers are key to understand where tipping is happening and we show that the

most dramatic racial changes are occurring at the boundaries of racial clusters. This pattern is predicted by our simple discrete choice demand system integrating Schelling’s bounded neighborhood and spatial proximity models. We interpret its predictive power for cross-sectional patterns and decadal changes in segregation as an argument in favor of modeling spatial spillovers when investigating the impacts of local policies that could lead to racial resorting. Especially within structural quantitative models of racial sorting in cities, capturing local residential spillovers could be key to answering policy-relevant questions.

Relation to the literature

Our paper adds to several literatures. In our theoretical model we build upon the fundamental work of Schelling (1969, 1971) and show how to nest his bounded neighborhood (including tipping) model inside a discrete choice framework that also captures the central forces of his spatial proximity or checkerboard model. Though our approach is simple, to the best of our knowledge the only other attempt of this integration has explicitly been made in Zhang (2011). The author follows the original checkerboard approach of Schelling and models neighborhood dynamics through an agent-based simulation of individuals with random taste draws moving between houses arranged on a grid.³ The model excludes other factors such as different price sensitivities across groups or housing supply. In contrast to Zhang (2011), we write a model of neighborhood choice that operates on a more aggregate level on which data is available more frequently. This allows for factors beyond the racial composition of a location and has the advantage that it can be taken to the data. We also move beyond Zhang (2011) in that we test predicted patterns of our model empirically using US census data.

Our model closely relates to a fast evolving literature in urban IO and quantitative spatial economics which models discrete neighborhood choice and endogenous sorting of heterogeneous groups. The literature following the IO tradition of demand estimation goes back to the static models of Bayer and Timmins (2005, 2007) and Bayer et al. (2007). A recent paper in this vein is Almagro et al. (2023) who use the HOPE VI public housing demolitions in Chicago to estimate racial preferences structurally and also investigate the welfare impacts of the policy. In the closely related quantitative spatial literature following traditions of trade economics, static models of neighborhood choice with multi-group sorting also exist. Examples in the context of income sorting include Tsivanidis (2023) or Couture et al. (2023) and an application to racial sorting is presented in Weiwu (2023). Going beyond the static tra-

³Other agent-based models of tipping are summarized in Axtell and Farmer (2022) and include Zhang (2004) or O’Sullivan (2009).

dition, dynamic models of neighborhood choice have been developed and estimated recently including Bayer et al. (2016); Caetano and Maheshri (2023) and Davis et al. (2023).

A common feature in all of these discrete choice models is that racial tastes are geographically constrained to depend only on the racial composition of the location under consideration. We take a different tack. We likewise develop a discrete choice model. But, crucially, we focus on a case in which racial tastes extend beyond an individual location. Evidence of very local preference spillovers is for example presented in Bayer et al. (2022). We also consider more distant spillovers in our model which could also be due to neighborhood networks as investigated in Ferreira et al. (2023). The paper most closely related in this regard is Bagagli (2023) who models racial spillovers in a fashion very similar to ours. She estimates them using quasi-exogenous variation in distances induced by the construction of expressways in Chicago during the 1950s. Rather than estimating the model in a specific local context, we instead develop simulations and derive empirical hypotheses that can be tested across all major US metropolitan areas. This procedure allows us to see more broadly whether the data is consistent with our hypotheses.

Our empirical investigation speaks to a recent paper by Dai and Schiff (2023) which examines ethnic clusters in US cities across multiple decades. Similar to their paper we also develop a methodology to identify clusters and examine price gradients within these. While their investigation distinguishes many ethnic groups and measures developments over time, our approach is restricted to a dichotomous world, focuses on race, and can explain through our structural model why certain patterns emerge. Our novel visualizations also add to the literature on measuring segregation and its secular trends (Massey and Denton, 1988; Glaeser and Vigdor, 2012).

Lastly, we contribute to the empirical literature on the estimation of racial neighborhood tipping points. We show that the reduced form approach advanced by Card et al. (2008) leads to erroneous results which are driven by biased White suburbanization.⁴ Although their idea to identify tipping points in a reduced form manner and then connect the magnitudes of these to different strengths of racial preferences in the cross section appears attractive at first, it proves susceptible to the confounding factor of White entry into suburban tracts.

We think that part of the reason why Card et al. (2008) come to their conclusions is the partial equilibrium approach they use to motivate their empirical analysis. While we do not

⁴Another relevant reduced-form paper on tipping is Easterly (2009). He investigates the same data as Card et al. (2008) but finds more “white flight” out of neighborhoods with a high initial share of whites than out of more racially mixed neighborhoods.

suggest an alternative reduced form strategy to identify tipping points, we instead argue for a structural approach to estimating racial preferences emphasizing spatial spillovers. In this regard, we are sympathetic to Caetano and Maheshri (2017), who estimate tipping points structurally in the context of school choice, and Blair (2023), who uses a methodology similar to Caetano and Maheshri (2017) to identify racial tipping points on the census tract level. Blair (2023) also investigates the role of outside options for tipping points. We think of these general equilibrium aspects as key for making sense of the relationship between tipping, White suburbanization, and the findings of Boustan (2010). We would however like to highlight the additional importance of racial clusters and spatial spillovers in the role of racial neighborhood change.

Outline

The remainder of the paper is structured as follows. Section 2 introduces our model of discrete neighborhood choice with spatial spillovers. It also examines how the model nests Schelling’s ideas of tipping, explains our simulation procedure, and introduces our main empirical hypotheses. In the following section 3 we assess our hypotheses empirically. In light of our findings, section 4 revisits results by Card et al. (2008), and section 5 concludes.

2 A Model of Neighborhood Choice with Spatial Spillovers

The theoretical ideas of Schelling (1971) are transparent, enormously influential, and yet difficult to take to data. For this reason, formal empirical work based on his models ranges from sparse (bounded neighborhood and tipping models) to non-existent (spatial proximity model). In this section, we will aim to show that many of Schelling’s main ideas can be introduced into a common framework that nests them in a discrete choice model, which allows us to create a set of predictions we can investigate with data.

The first step towards arriving at a set of testable predictions from his models is to write down the general version of the discrete choice model with which we will work. Following the existing discrete choice work on racial preferences over locations, we allow that there are idiosyncratic tastes that affect the value of locations to individuals in addition to the local racial composition. In contrast to this existing discrete choice work, our general model allows for individuals to care about the racial composition not only of their own but also other nearby locations.

With this general version of the discrete choice model in hand, we show how it nests the traditional bounded neighborhood model with tipping as a special case. The key assump-

tion is to consider the model in a strict partial equilibrium setting. We then go on and create a bridge to the full spatial proximity or checkerboard model by allowing for many locations among which agents can choose. But in this step we maintain the assumption that individuals care only about the racial composition of their own location. This permits us to introduce adding-up constraints across multiple locations into the bounded neighborhood model. Finally we develop a version of the full spatial proximity model, which maintains multiple locations but also allows individuals to care about the racial composition not only of their own but also of proximate locations. Introducing these in stages makes clear the empirical predictions that emerge from each stage.

We now describe a discrete choice model that incorporates main ideas from Schelling’s spatial proximity model and can also generate tipping behavior as in the bounded neighborhood model.

Geography Space consists of a discrete set of locations $j \in J$ endowed with a distance metric d_{jk} that describes the distance between two locations j and k . In our empirical analyses we will focus on census tracts.

Demand There are different population groups $r \in R$ living in the city each having an exogenous total size of N_r . The total population inhabiting the city is thus $N = \sum_{r \in R} N_r$. Following a simple logit specification, households i of group $r(i)$ derive the following indirect utility from living in j

$$v_{ij} = u_{r(i)j} + \epsilon_{ji} \tag{1}$$

where ϵ_{ji} is a household- and location-specific i.i.d. Gumbel shock with unit scale. The location specific mean utility is common across groups and takes the following form

$$u_{r(i)j} = -\alpha_{r(i)} \log(p_j) + \beta'_{r(i)} \sum_{k \in J} w_{jk} s_k + \eta_{r(i)j}. \tag{2}$$

Here, p_j is the average price at location j and $s'_k = (s_{1k}, \dots, s_{Rk})$ is a vector of neighborhood group shares at location k . Neighborhood fundamentals and amenities which are not affected by racial sorting but which can be group-specific are captured through η_{rj} . α_r describes the price sensitivity and β_r captures racial preferences of group r . In the baseline formulation of our model racial preferences enter indirect utility linearly but the setup can principally be extended to more complex cases.⁵ Spillovers across neighborhoods depend on the origin-

⁵Note that this indirect utility formulation can also be derived by assuming a Cobb Douglas utility function with a housing share of $\tilde{\alpha}_r$ and a multiplicative taste draw that is i.i.d. Fréchet distributed across locations with shape parameter $1/\sigma_r$, location parameter 0, and scale parameter 1. The log indirect utility

destination-specific weights w_{jk} . They describe the degree to which households living at j care about the racial composition in location k . In principle, the distance decay could have an arbitrary form varying by race and depending on population density or the local transportation network. In the initial version of our model, we assume that w_{jk} decays exponentially with distance, its decay rate is determined by κ , and weights satisfies the normalization $\sum_k w_{jk} = 1$:

$$w_{jk}(\kappa) = \frac{e^{-\kappa d_{jk}}}{\sum_{k' \in J} e^{-\kappa d_{jk'}}$$

Using this formulation, racial preferences are very localized if $\kappa \rightarrow \infty$, meaning that households only care about the racial composition at location j itself. By contrast, only the city-wide racial composition plays a role if $\kappa = 0$. The incorporation of these spatial spillovers bridges the gap between Schelling's spatial proximity model and most existing discrete choice models of segregation.

Households choose where to live by maximizing their utility. This yields the following aggregate demand of group r for location j depending on the vector of all prices, neighborhood racial shares, and exogenous demand shifters:

$$D_{rj}(\{p_k\}, \{s_k\}, \{\eta_{rk}\}) = N_r \frac{\exp(u_{rj})}{\sum_{k \in J} \exp(u_{rk})} \quad (3)$$

Supply In our baseline scenarios, supply for housing is fixed and exogenous. Each location is endowed with a fixed housing stock H_j and total housing units available equal the total population: $\sum_{j \in J} H_j = N$.⁶

Equilibrium The share s_{rj} of group r at location j is given by

$$s_{rj} = \frac{D_{rj}(\{p_k\}, \{s_k\}, \{\eta_{rk}\})}{\sum_{r'} D_{r'j}(\{p_k\}, \{s_k\}, \{\eta_{rk}\})} \quad (4)$$

An equilibrium is then defined by a vector of prices $\{p_k\}$ and a matrix of racial shares $\{s_{rk}\}$ such that equation 4 is satisfied and housing supply equals total demand at each location

$$H_j = \sum_r D_{rj}(\{p_k\}, \{s_k\}, \{\eta_{rk}\}). \quad (5)$$

in this formulation would be equivalent to our additive formulation in equations (1) and (2) with our price sensitivity α_r corresponding to the housing share multiplied with the shape parameter $\alpha_r = \tilde{\alpha}_r / \sigma_r$.

⁶The model can easily be extended to allow for endogenous supply for example by specifying a simple reduced form housing supply curve such as $H_j = \bar{H}_j p_j^{\theta_j}$ where \bar{H}_j is a supply shifter and θ_j is the local supply elasticity.

The existence of an equilibrium can be shown through Brouwer’s fixed point theorem but uniqueness is not guaranteed (Bayer and Timmins, 2005). In fact, multiple equilibria are a key feature of the model and arise if agglomeration forces in the form of racial preferences and spillovers are sufficiently strong.

In the following sections, we will show how our framework can nest variants of Schelling’s bounded neighborhood and spatial proximity models. We will also use variously restricted models as a basis for simulations from which we will develop hypotheses to be examined empirically in Section 3. In all of these sections we will focus on two racial groups $R = \{w, m\}$ where w indicates White and m indicates Minority households⁷. Since racial preferences enter linearly in the baseline version of our model, we can describe a location’s racial composition by its scalar Minority share s_{mk} . We limit ourselves to this simplified two-group analysis, as much of the existing literature on tipping is written in a two-group context and the simplified setting allows us to gain intuition more easily.⁸ Most of what follows can be extended to a multi-group setting.

2.1 The Bounded Neighborhood Model in Partial Equilibrium and Tipping

The model outlined previously closely relates to the bounded neighborhood model and ideas about tipping, as shown through the lens of bid rent functions in Becker and Murphy (2000) and Card et al. (2008). To see this relationship more clearly, we move from the more general equilibrium world that simultaneously describes demand across all locations to a partial equilibrium setting that focuses at a single location j . If we assume that demand (and thus prices, racial shares, and fundamentals) at all locations outside of j does not respond to changes in racial shares at j , we can derive j -specific partial equilibrium bid rent functions. While such an approach ignores adding-up constraints and spillovers, it corresponds precisely to Schelling’s bounded neighborhood approach.

A partial equilibrium bid rent-function $b_{rj}(s_{mj})$ describes the maximum willingness to pay of a marginal household of group r to move into location j if the Minority share at that location is s_{mj} . It is implicitly defined through the equation

$$D_{rj}(b_{rj}(s_{mj}), s_{mj}, \eta_{rj}) - H_j s_{rj} = 0,$$

⁷We take non-Hispanic White individuals to be “White”, and all other races to be “Minority”. The two groups defined here, in our theoretical portion, align with definitions of racial groups we use for analysis in our empirical section.

⁸For example Schelling (1969, 1971); Becker and Murphy (2000); Card et al. (2008) and Easterly (2009) all consider a two group context. Only more recently Caetano and Maheshri (2017) have explored multi-group tipping in a school choice setting.




Figure 1: Illustration of bid rent curves under different homophily preferences. Minority bid rent curves are colored red. White bid rent curves are colored blue. Only White’s homophily preferences are varying across the different illustrations. Stable equilibria are indicated by a filled black circle. In the case of strong homophily preferences, multiple equilibria emerge. The unstable equilibrium is indicated by a non-filled circle and tipping points are indicated by dashed vertical lines.

assuming that $\{p_k, s_{mk}, \eta_{rk}\}$ remain unchanged for $j \neq k$. Figure A1 in the appendix provides graphical intuition for how the bid rent curve is constructed. As both, the White and the Minority bid rent curve depend on s_{mj} , they can be plotted in the same diagram with crossings pinning down (partial) equilibrium Minority shares and neighborhood prices. Stable equilibria are characterized through $b'_{wj} < b'_{mj}$ while intersections where $b'_{wj} > b'_{mj}$ are unstable equilibria.⁹ Changes in population sizes N_r or location fundamentals η_{rj} (relative to outside options) can shift the bid-rent curves up and down. If such shifts lead to an intersection where $b'_{wj} = b'_{mj}$ this is a tipping point.

Figure 1 illustrates bid rent curves for some selected cases. White bid rent curves are colored blue while Minority bid rent curves are displayed in red. Panel (a) illustrates a setting with weak homophily preferences of Whites, no racial preferences of Minorities, and a unique partial equilibrium. Panel (b) shows a setting with strong homophily preferences of Whites in which multiple equilibria exist. The difference in shape of the bid rent functions across the two panels emphasizes a characteristic feature of homophily in the model, namely that it makes the bid rent function more elastic. For Minority households, who are assumed to have no homophily preference here, the bid rent function has the conventional downward sloping shape. For White households, the bid rent function is more elastic in panel (a) and even

⁹Caetano and Maheshri (2017) refer to such unstable equilibria as tipping points using a closely related “S” shape approach to identify them. We follow the approach of Card et al. (2008) in what we label a tipping point in partial equilibrium.

becomes upward sloping in panel (b) due to homophily preferences. In both cases, shifts in either of the two bid rent functions can lead to drastic changes in equilibrium Minority shares. With weak homophily preferences, adjustments occur without involving a bifurcation and crossing of a proper tipping point as there is no Minority share for which $b'_{wj} = b'_{mj}$. In (b), an upward shift of the Minority bid rent curve can render the low Minority share equilibrium unstable leading to a bifurcation. The tipping point where $b'_{wj} = b'_{mj}$ is indicated by the left vertical dashed line. If the White bid rent function shifts upwards, an analogous tipping point exists at the right vertical dashed line.¹⁰ The fact that drastic racial change in the bounded neighborhood model is possible both when a formal bifurcation exists (panel b) and when it does not (panel a) is a caution that drastic change alone is not evidence of such bifurcations. It also suggests the potential value of considering such drastic change directly as tipping without only or necessarily focusing on a search for bifurcations.

The bounded neighborhood model provides a rich setting for contemplating the fate of a single neighborhood in response to shocks. However it also lays bare a shortcoming of the framework, particularly when we turn to understanding the evolution of all census tracts within an MSA. The aggregate shock contemplated, that of an upward movement in the Minority bid rent curve, in all three of these cases leads to White *exit* from the neighborhood. This is true whether or not tipping is formally present. When we try to apply this to understanding what happens for all tracts in an MSA, it is clear that one must move to general equilibrium in order to understand the adding-up constraints that affect White *entry* to other neighborhoods. Everyone must go somewhere.¹¹ This is discussed further in Section 2.4.

2.2 Foundations for Simulations of the Discrete Choice Model

Having explained how our discrete choice model relates to bid rent functions and how it can give rise to tipping, we return to the analysis of the aggregate demand system in the following. We want to understand the general equilibrium properties of the model - both when racial preference spillovers across neighborhoods are turned off (i.e. $\kappa \rightarrow \infty$) and when they exist ($\kappa < \infty$). We accomplish this through a number of simulations that lead us to

¹⁰Our approach to the bounded neighborhood model has similarities and contrasts with the model developed by Card et al. (2008). The model does admit the possibility of a tipping point at a critical Minority share s_{mj}^* . However one feature of the logit demand model relevant here is that the bid rent curve of group r will always tend to infinity for low values of s_{mj} . This ensures that in every location households from every group are represented in expectation. In the post-tipping equilibrium, the Minority share won't be strictly $s_{mj} = 1$, which is in any case rare in the data.

¹¹This will prove important when we seek to understand the roles, respectively, of White exit and White entry in driving prominent empirical results in the bounded neighborhood model in Card et al. (2008).

several empirical hypotheses, listed in section 2.4. We examine these hypothesis in sections 3 and 4.

To ease visual exposition we consider a city with $J = 100$ neighborhoods that are located on a unit square in a 10×10 grid throughout all simulations. Each location has the same exogenous housing supply $H_j = 1$ and each group has the same unobservable preference shock for every location $\eta_{rj} = 0$, i.e. all locations are identical in fundamentals.¹²

We simulate cities with two population groups: Minorities m and Whites w . Each group has equal price sensitivity $\alpha_m = \alpha_w = 20$ and Whites have a strong preference to co-locate $\beta_w = 8$ while Minorities are indifferent about the racial composition of their neighborhood $\beta_m = 0$. The chosen parameters imply a semi-elasticity of $\beta/\alpha = 8/20 = 0.4$, i.e. if the average Minority share in the neighborhood increases by 1 percentage point, prices must decrease by 0.4% to keep White households indifferent. White households are in the majority with $N_w = 70$ while the Minority constitutes the remaining $N_m = 30$.

We break the simulations into two analytically distinct cases. The first is a limiting case of the general model, which places the bounded neighborhood model in general equilibrium and considers outcomes when the influence of proximate neighborhoods is vanishingly small (i.e. $\kappa \rightarrow \infty$). The second case is our general case of the spatial proximity model, where there is a clear effect of the proximate neighborhoods' racial characteristics on the utility of agents ($0 < \kappa = 15 < \infty$). Together with our assumption of locations on a unit square, $\kappa = 15$ implies that the racial composition of the considered tract contributes on average about 45%, the 8 neighboring tracts together contribute another 45%, and all other tracts the remaining 10% to the experienced racial composition at the location.¹³

To find an equilibrium, we initialize each location j with a minority share s_j^0 that is independently drawn from a uniform distribution. Given the initial minority shares $\{s_j^0\}$ we find the price vector $\{p_j^0\}$ that balances supply and demand at each location. Given this price vector, we can again update minority shares. We iterate between updating prices and minority shares until we converge to an equilibrium satisfying equations 4 and 5.¹⁴

To summarize and visualize key features of the simulated equilibria we proceed in the following manner. First, we find clusters of contiguous same-race-mode locations. Second,

¹²Simulations can easily be extended to a setting with many more locations, endogenous housing supply, and heterogeneity in fundamentals. We programmed the model in Julia and with 2500 locations and multiple groups, the model usually still converges within seconds.

¹³The implied spatial weights $w_{jk}(\kappa)$ for the two cases are illustrated in appendix figure A2.

¹⁴This procedure for finding the static equilibrium mimics a myopic dynamic adjustment of households who move without anticipating the impact on the local racial composition.

we assign an integer distance to each location that describes how many tracts one has to traverse to reach the border between clusters of opposite mode. We assign negative numbers to locations that are inside of a Minority cluster and positive numbers to locations inside a White cluster. For example, a White-mode tract bordering a Minority tract is assigned the location 1, those further inside the cluster being assigned in turn $\{2, 3, \dots\}$. Third, we provide diagrams that show the median fraction Minority of all tracts with the same distance from a cluster boundary.

When we later turn to the empirical investigation using actual census tract data we use the same methods of clustering and slicing through the data to summarize key features of cities' racial segregation. As we will see, this representation of the data converts the multiple equilibria and hysteresis strongly present in the data into features that we care about, that are robust across different equilibria, and that allow meaningful comparisons between simulations and data.

2.3 The Bounded Neighborhood Model in Simulated General Equilibrium

We begin with a special case that acts as a transition between the bounded neighborhood model and our full spatial proximity model. This places the bounded neighborhood in general equilibrium, so that adding up constraints now bind. However, we maintain the key feature of the bounded neighborhood model that outcomes depend directly on the racial composition of one's own neighborhood and no other. Most of the existing empirical discrete neighborhood choice literature on neighborhood sorting follows this approach, abstracting from preference spillovers.¹⁵ Limitations that we discuss in the following thus also apply to most of the existing models.

As one would expect, the resulting clustering in the left panel of Figure 2 is modest, reflecting random variation. In this example, no Minority neighborhood exists that does not also border a White neighborhood. In the right panel, one sees that even in 100 randomly initialized simulations of this bounded neighborhood model, the depth of the Minority cluster never exceeds -2 , i.e. one layer of depth beyond contact with the White cluster.¹⁶ In addition, as we move from the edge of the Minority cluster (position -1) inward or outward into the White cluster, there is no variation in the Minority share within cluster types. This is exactly as one would expect given the absence of spatial racial spillovers. This thus also gives rise

¹⁵A notable exception is Bagagli (2023) though she does not focus on the importance of racial clusters that we emphasize.

¹⁶Even this fails to make clear the rarity of Minority clusters in the simulations. Across 100 simulations, 75 had zero locations of -2 or further to the interior; 22 had exactly 1 such location; and only 3 had two such locations.




Figure 2: Simulation of the Bounded Neighborhood Model in General Equilibrium. The left panel is a typical single equilibrium. The right panel shows an average of Minority share by location for 100 simulations. Absent spillovers across locations, the clustering of Minority tracts is weak.

to a very steep gradient at the cluster boundary. These features will figure importantly in our discussion of the data and the importance of spatial racial spillovers.

2.4 The Spatial Proximity Model in Simulated General Equilibrium

We now develop the general version of our spatial proximity model. In addition to placing the model in general equilibrium, we now allow individuals to have preferences not just for the racial composition of one's own neighborhood but also for the racial composition of neighbors. This is accomplished in our model by setting the spatial preferences decay to $\kappa = 15$ putting weight on one's own neighborhood as well as nearby tracts.

Figure 3 shows the results of a baseline simulation. Panel (a) shows for this iteration the Minority share for each square in the checkerboard. If we define a cluster as the set of contiguous tracts with the same mode, then this example shows one cluster for each group. Moreover, there are no isolated squares or even mini-groups of isolated squares within the other cluster. One-hundred percent of the population lives in one or the other large cluster.

Panel (b) provides a simplified view of the resulting geography that orders locations according to their distance from the Minority cluster boundary with the White cluster (so location -1 is just inside the Minority cluster). The panel shows an average for 100 simulations. Two features stand out clearly. The first is that the locations strictly interior to each of the clusters have extreme degrees of segregation, with typical values of the Minority share at 98% interior to the Minority cluster and 2-3% interior to the White cluster. This is true in spite of the simulations assuming complete indifference by Minorities about the racial composition of their locations. Second, at the boundaries of clusters, there is a very

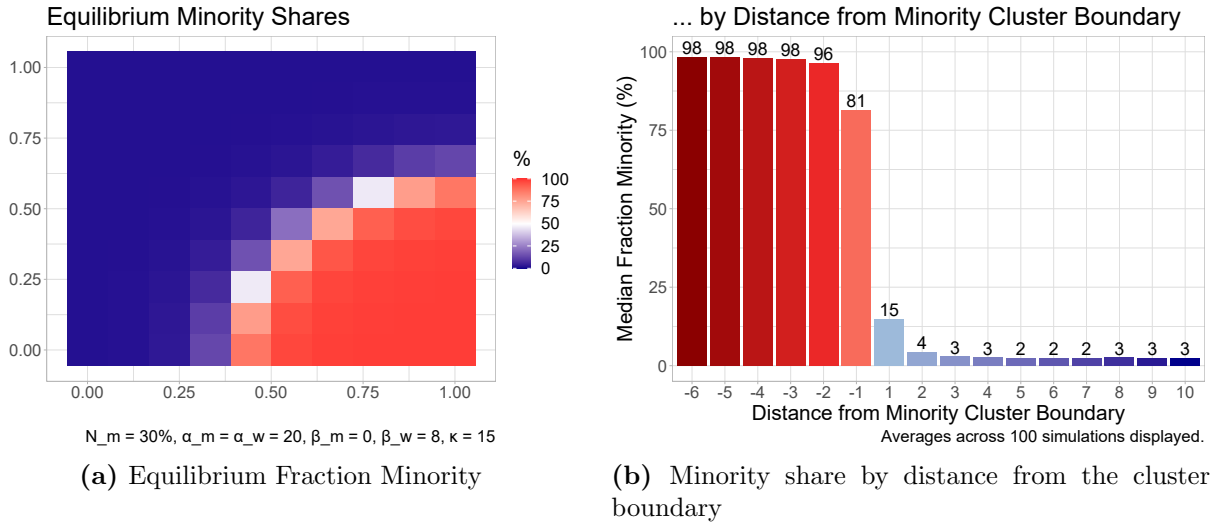


Figure 3: Simulation of the Spatial Proximity Model in General Equilibrium. The left panel is a typical single equilibrium. The right panel shows an average of Minority share by location for 100 simulations. In the presence of racial spillovers across locations, the clustering of Minority tracts is strong.

sharp racial gradient, yet it is softened by the location near the edge of the boundary. The logic is that some Whites with particularly strong Gumbel shocks for these boundary neighborhoods are willing to locate there (but many fewer further interior to the Minority cluster) because of the proximity to the high White share locations, given the spatial decay of racial preferences for Whites. This can be thought of as a "market access" effect in terms of racial preferences.

Observing outcomes from the simulation we form two initial hypothesis that we will test in the data.

Hypothesis 1 (Existence of Schelling Clusters). When racial preferences are important, a large fraction of tracts will be found in contiguous same-mode clusters. These are akin to the segregated neighborhoods in Schelling’s spatial proximity, or checkerboard, model.

Hypothesis 2 (Strong Segregation Interior to Clusters:). A strong prediction of Schelling-type model is that genuinely racially mixed neighborhoods don’t survive. In our variant of the bounded neighborhood model, outcomes tend toward strong but not perfect segregation, due either to formal tipping or because of highly elastic bid rent functions deriving from homophily. In the spatial proximity model, individuals tend to flee locations surrounded by the other group, and so especially the interior of the respective clusters will feature quite high levels of segregation.





Figure 4: Equilibrium Prices in the Spatial Proximity Model. Those with strong racial preferences pay a premium, even absent other sources of variation.

Price Gradients A notable feature of the equilibrium is that those with strong racial preferences pay a price (cf. Cutler et al., 1999). This is illustrated in Figure 4. The simulation is perfectly symmetric across groups and locations except for the strong homophily of Whites and absence of homophily for Minorities. For example, no income differences have been introduced. Averaging across 100 simulations, the median price deep in a White cluster is 20% higher than its counterpart in a Minority cluster. There is virtually no price variation inside the Minority cluster, reflecting the asymmetry in willingness to pay to be remote from Whites. The price gradient inside the White cluster near the clusters’ boundary is strong, but not precipitous. This leads us to another hypothesis.

Hypothesis 3 (Price Gradients Inside White vs. Minority Clusters). When Whites have strong homophily and Minorities have none, location prices will rise strongly in movements away from the boundary inside the White cluster, but will have little variation inside the Minority Cluster.

Shocks to the Minority Share One of the critical questions in the segregation literature is how neighborhoods evolve inside the city in response to a shock to the aggregate city Minority share. Much of this is motivated by the shocks of the two great migrations of Blacks to northern cities. Since our model features multiple equilibria, one needs to start with an initial equilibrium and then consider what happens relative to this baseline as we progressively raise the Minority share. We initialize the city with the original equilibrium Minority shares and then increase the citywide Minority share while simultaneously reducing the White share. This is depicted in Figure 5 as the Minority share in the city rises in steps from 0.3 to 0.5.

Adjustment to Changes in City Minority Share



Change relative to baseline




Figure 5: Increases in the Minority share lead to racial change at the boundary of clusters.

The first key insight from the exercise is that the clusters remain intact, only adjusting their sizes. The lower panel of Figure 5 shows this in changes, where these are strongly concentrated at the boundary of the clusters. These features are very much to be expected from a spatial proximity model, since it is the same forces that lead to the initial sorting into clusters that preserves the clusters under change. The observed patterns of change form the basis for two more hypotheses.

Hypothesis 4 (Change Occurs at the Boundary of Clusters). When there is a shock to the racial composition of the city, change in racial shares is focused on the boundaries of the racial clusters. The logic of segregation that gives rise to clusters in the first place asserts itself to preserve the clusters, insuring that change is most likely at the boundary of clusters.

Hypothesis 5 (Tipping as Drastic Change). When we interpret tipping as drastic change in a period, tipping will be focused on the boundaries of clusters, and the greater the threshold of change required to qualify as tipping, the more strongly change will be focused on the boundary. This follows from the combination of a strong gradient at the boundary of clusters and that change occurs at the boundary of clusters.

Biased White Suburbanization The last exercise highlighted the important feature that, from an initial equilibrium, changes in race shares lead to change primarily at the

boundaries of clusters. This suggests the value of an exercise that incorporates additional features of the American context. One of these relates to the fact that historically many Minority mode areas were referred to as the “inner city,” since they were very frequently in areas central to the city. The second is that while Minority presence did increase in MSAs, the total White population didn’t typically decline, so a focus on shares misses part of the story. Third, the areas initially suburbanized or at the suburban frontier very often have a White mode. This articulation of the setting is very sympathetic to the approach of Boustan (2010).

It also allows us to articulate a competing hypothesis to consider when reviewing evidence advanced in favor of tipping. Call it the *biased White suburbanization hypothesis*. Tipping theory is premised on the idea that a tract being above or below a common tipping point s_{jm}^* suffices to identify neighborhoods that do or don’t tip. Excluding neighborhoods that experience explosive *growth* shouldn’t eliminate the contrast in experience. Moreover, the biased White suburbanization hypothesis, in contrast to tipping, specifies the location of tracts where great entry will occur.

To put it starkly, in tipping theory the contrast in experience between s_{jm} below or above s_{jm}^* is about differential rates of *exit*, while the White suburbanization hypothesis says that the critical element is driven by drastic *entry* to $s_{jm} < s_{jm}^*$ tracts remote from Minority clusters.

Hypothesis 6 (Biased White Suburbanization). If, as in many U.S. cities, the Minority cluster is initially situated in a central “inner city,” then growth in that city’s Minority population will lead to an expansion of the central Minority cluster and growth in the White population in low Minority ($m < m^*$) locations at a suburban edge.

3 Evidence on the Spatial Proximity Model from a Panel of Cities

In this section, we examine the hypotheses developed in the preceding section empirically using U.S. Census data. We focus on the time period from 1970 to 2000, as we have detailed information on the racial composition of census tracts as well as a rich set of covariates available for these four censuses. Our panel data of cities derives from the replication data of Card et al. (2008). It does not go beyond the year 2000 as we want to compare the results that we will show to their prominent investigation of tipping points. All data is harmonized to 2000 census tract geometries and, similarly to Card et al. (2008), we focus on tracts that

are located within 1999 MSA definitions.¹⁷ As in the preceding theory section, we focus on a dichotomous classification of race comparing non-Hispanic White households and minority households, understood to be the complement. For simplicity and following the literature, we refer to the two groups as White and Minority households respectively.¹⁸

3.1 Schelling Racial Clusters in the Data

Schelling’s spatial proximity, or checkerboard, model is noted for its emergent property of strong spatial segregation. However, this is a model of individuals, while our discrete choice model with spatial racial spillovers focuses on the composition of neighborhoods and our data concerns census tracts. To verify the relevance of the spatial proximity model in the data, we need to operationalize the idea of Schelling racial clusters.

To operationalize the model, we classify census tracts according to their racial composition into White and Minority mode locations. The classification procedure is illustrated in panels (a) and (b) of Figure 7 for selected census tracts in Chicago’s South Side. We then identify racial clusters as a set of contiguous tracts within an MSA with the same modal race. One has to make a choice of the minimum number of contiguous tracts that will constitute a cluster. We illustrate this for varying cutoffs in Figure 6. Keeping in mind that a typical-sized tract will have roughly 4,000 residents, a cluster of 20 tracts requires roughly 80,000 people living in contiguous same mode tracts. At this cutoff, and in all years we observe, between 77% and 87% of all Americans live in own-race racial clusters. This number rises to above 80% in all decades if we consider a minimum cluster size of 5 tracts.¹⁹ In short, Schelling racial clusters are a first order feature of American cities, consistent with hypothesis 1.

3.2 Clusters, Location, and Degrees of Segregation in the Data

Our discrete choice model with spatial spillovers suggests, per hypothesis 2, that there will be strong segregation by race interior to clusters, but that this will be attenuated at cluster boundaries. To investigate this hypothesis empirically we slice through the census data in the same manner that we used to display our simulation results. We identify cluster boundaries

¹⁷We also follow the sample selection methods put forward in Card et al. (2008) and exclude all tracts in which (1) the decadal population growth rate exceeds the MSA mean by more than five standard deviations, (2) the ten-year growth in the white population exceeds 500% of the base-year total population, (3) the MSA contains fewer than 100 tracts (after applying the previous criteria).

¹⁸We also note that during the time period considered, non-Hispanic White households were constituting the majority population in our sample.

¹⁹There is variation across groups. Consistent with their majority status, Whites are more clustered, have a larger share of their total population living within their own clusters, and are declining over time in each dimension. Minorities are less clustered, have a smaller share of their total population living within their own clusters, but are rising over time in each dimension.



Figure 6: Racial clusters dominate US cities. Even with high thresholds for the number of contiguous tracts required for inclusion, the vast majority of Americans live in own-race racial clusters.

in each MSA and then classify tracts according to their distance from the boundary. This is illustrated in panel (c) of Figure 7. The boundary is treated as the origin with distances $\{1, 2, \dots\}$ assigned to White mode tracts and distances $\{-1, -2, \dots\}$ for the respective Minority tracts. Tracts further to the interior of Schelling racial clusters have darker shades. Panel (d) projects these locations onto a simpler space displaying distance to the boundary of racial clusters on the horizontal and median Minority shares on the vertical axis.

Having outlined the methodology, we begin our examination with the case of Chicago in 1970, then extend this to all decades, and later to all MSAs.

Segregation by Location in Chicago

Figure 8 shows the bar graph for locations inside the Chicago MSA in 1970 through 2000.²⁰ Two patterns in the 1970 data become clear. First is that the tracts strictly interior to the Minority cluster are quite highly segregated, with these tracts having a median Minority share of 0.98. Even the tracts at location (-1), just inside the Minority cluster, have a median $s_{jm} = 0.91$. Moreover, as soon as we pass into the first set of White tracts outside the Minority cluster, there is a dramatic fall in the median Minority share to $s_{jm} = 0.13$. This fall especially stands out as the way in which we classified tracts into White and Minority mode does not force such a pattern to appear. Across all White tracts, the median Minority share is only $s_{jm} = 0.03$.

²⁰Figure A3 in the appendix provides the corresponding map of Chicago in 1970 for the MSA and the central city.

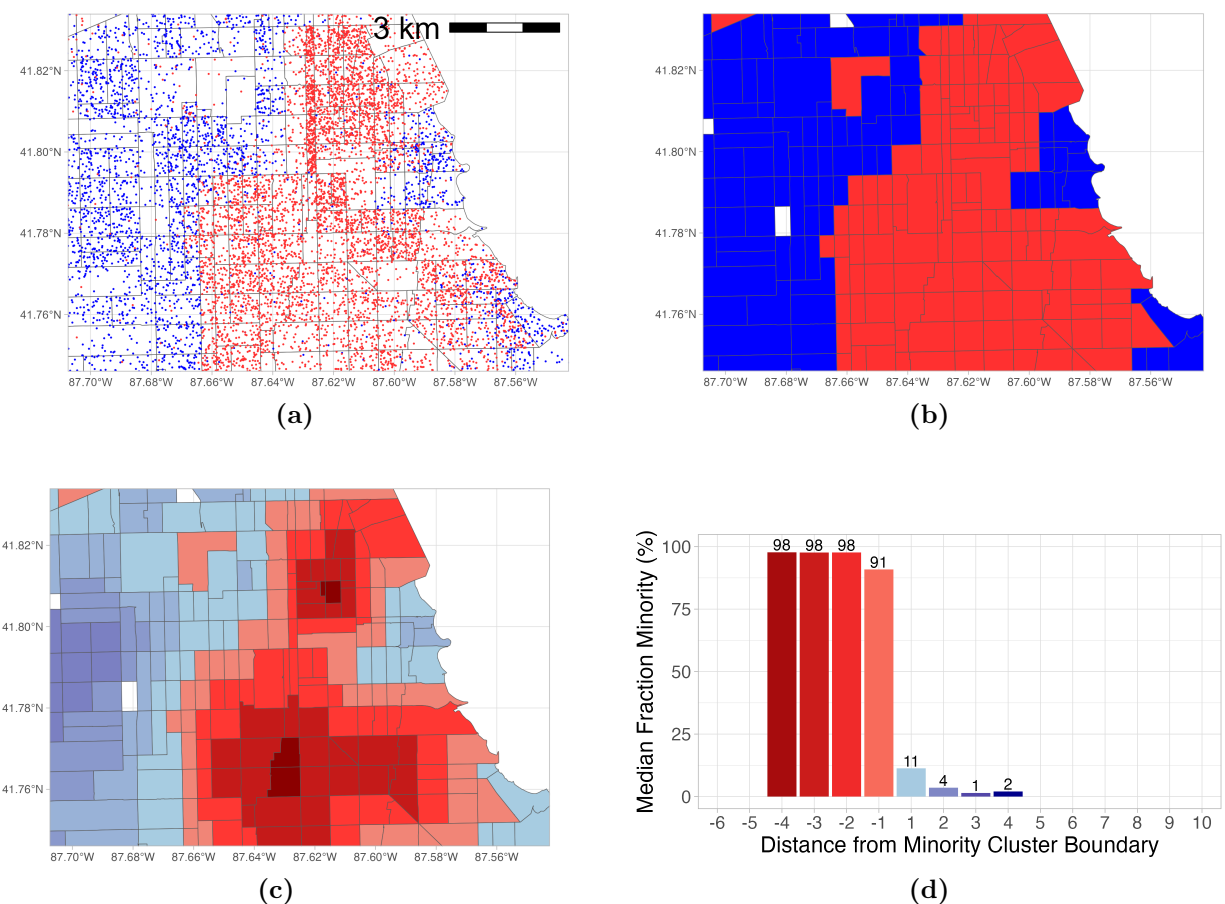


Figure 7: Minority and White population illustrations for the South Side of Chicago, 1970-1980. We show how to move from tract population data on race to our representation of minority share by location. Panel (a) illustrates population counts by census tracts; each dot represents 100 individuals, with red dots representing Minorities and blue dots representing Whites. Panel (b) colors census tracts by modal race given the underlying data from panel (a). Panel (c) shades tracts by distance from Minority cluster boundary, with lighter colors indicating closer to a boundary. Panel (d) is a histogram of median fraction Minority by distance bin for all tracts in this sample. Census tract geometries are 2000 geographies; cluster size is set to 1.

We first look to see if the Chicago 1970 pattern continues to hold in the following decades in Figure 8. Several important insights are revealed. The first is that Minority clusters, particularly the interior, are hyper-segregated with values close to the theoretical prediction of perfect segregation. Second, over time there does seem to be some diminution of the gradient at the boundary of Minority and White clusters. Third, there is an expansion of the count of interior layers in the Minority clusters from 4 to 6, suggesting their geographical expansion. Finally, there is a steady rise in the Minority share in the tracts far from Minority clusters, rising from a typical 3% in 1970 to circa 14% in 2000.




Figure 8: Median Minority Share by Location Across Decades, Chicago 1970-2000. Gradients are large at the boundary of clusters in Chicago, and while the interior of Minority clusters remain highly segregated across decades, the interior of White clusters becomes less segregated.

Segregation by Location Across All MSAs

The evidence from Chicago supports Hypothesis 2 on strong segregation inside clusters as well as the expected boundary gradient. To investigate if the Chicago patterns hold more generally, we look at evidence from all MSAs in Figure 9. Its panels strongly mirror the evolution that we saw in Chicago over the same period. As Glaeser and Vigdor (2012) document, 1970 represents a peak period of segregation in US MSAs as measured by the dissimilarity index, and they argue that the evolution represents the “end of the segregated century.” By the scalar dissimilarity index, this is absolutely correct. But Figure 9 underscores that the advances are uneven. The Minority mode clusters remain virtually unchanged in the interior, and indeed expand in size from 5 layers to 6. Most of the change appears to be the rise in the Minority share in White mode clusters. This is progress, but only partial.

3.3 Price Gradients by Distance from Cluster Boundaries

When White households have homophily preferences and Minority households have weak or no preference to co-locate, the spatial proximity model leads to a clear prediction regarding the gradient of neighborhood rental prices: they should be relatively flat within Minority clusters but rise the further one moves into a White cluster. In other words, White households




Figure 9: Median Minority Share by Location, All MSAs 1970-2000. Gradients are large at the boundary of clusters across all MSAs, and while the interior of Minority clusters remain highly segregated across decades, the interior of White clusters becomes less segregated, another fact robust across MSAs.

pay a price to segregate. As suggested through our simulation, this pattern should arise even in the absence of other neighborhood fundamentals that vary across space.

Empirically, we can investigate Hypothesis 3 using tabulated rental and house price data from the Census. This data is not available in the Card et al. (2008) replication files, so we added it using area and population weighting to adjust for changing census boundaries over time. Our investigation of prices is limited to the Chicago MSA only. We will expand it to other MSAs in the future.

Figure 10 shows median rental prices in the Chicago MSA from 1970 to 2000. To make prices easily comparable across decades, they are normalized by the rent paid in tracts that are 2 tracts away from the Minority cluster boundary. The observed price gradients support the main predictions of the model: within Minority clusters, gradients are relatively flat, while moving into the White clusters prices first increase substantially and then flatten out. Median relative house values displayed in appendix figure A5 show a similar pattern.

Three deviations from the predictions appear: (1) Especially in 1990 but also in 1980, there is a slight increase in rental prices as one moves deeply into the Minority cluster. This could be evidence of Minority homophily preferences. (2) In 1990 and especially in 2000, rental prices




Figure 10: Median Rental Prices relative to Prices in Tracts with a Distance from the Minority Cluster Boundary of -2, Chicago MSA. Prices are relatively flat inside Minority clusters and rise strongly as we leave the racial cluster boundary and enter the interior of White clusters.

far inside the White cluster are falling with increasing distance from the minority cluster. This could be the result of selection. With the increasing size of Minority clusters over the decades, tracts that are far away from the cluster boundary are more likely tracts that are very far outside of the city commanding lower rents. (3) In 1970, rental prices in tracts just inside the White cluster at a distance of 1 are below prices at -1. This could be a result of anticipation of tipping as Whites departed and Minority clusters were expanding rapidly during the 1970s. It could also be an indication of still-segmented markets resulting from strong racial discrimination, preventing Minority households from moving into relatively affordable housing just outside of the cluster edge.

Despite these deviations, the overall the observed price patterns are surprisingly well in line with the predictions of the simulation.

3.4 Decadal Racial Changes in Chicago and all MSAs

A central prediction that we obtained from the spatial proximity model is that change is most likely happening at the boundaries of clusters, where the groups are most directly in contact. The bounded neighborhood model suggests that when change happens, it will be drastic, so we focus on the location and magnitude of change as formulated in Hypotheses 4

and 5.

Racial Change in Chicago

We again start our inquiry by returning to the example of Chicago, focusing on the decade from 1970 to 1980. During this time period, more than 90% of Chicago tracts experienced a decline in White share. This included 30% who lost 10 percentage points (pp) or more; 15% that lost 25 pp or more; 7% that lost 50 pp or more; and some that lost as much as 98 pp. That this transpired over the course of a single decade indicates that drastic racial and ethnic change was common in Chicago neighborhoods in the 1970s.²¹

Figure 11 displays where change was happening in relation to clusters. It indicates that change is very strongly concentrated at the boundary of clusters. More than 40 percent of tracts that experienced a decline of 25 percentage points or more in the White population were in White mode tracts immediately abutting Minority mode tracts, and this figure rises to 60 percent if we add tracts in location 2 as we move away from the boundary with the Minority cluster. This is a powerful concentration of drastic change immediately adjacent to the Minority cluster.

We can illustrate this for different magnitudes of decline in the White share, from 10+ p.p. to 75+ p.p. The figure shows drastic change in this period is always centered on the White tracts immediately adjacent to the border of the Minority cluster, and the more drastic the change considered, the more strongly it is concentrated there.

Figure 12 allows us to visualize this spatially by focusing on the location of declines in White share of 25 p.p. or more for the Chicago MSA as a whole (left) and for the central portions of the MSA (right) where Minority clusters are principally located. The map powerfully reinforces this central message that drastic change is concentrated near the boundary of Minority clusters.

Racial Change in All MSAs

We can also now look at the correlate histograms pooling all MSAs in our sample for 1970-1980. These make clear that in this dimension, the case of Chicago is not an outlier. In the full sample, nearly 40 percent of the tracts with 25 percentage points or more are in the White tracts immediately bounding Minority clusters, and this rises to 60 percent if we include tracts in location 2 away from the Minority cluster.

²¹Figure A4 in the appendix provides a histogram of percentage point White share changes by Tract for Chicago in 1970. A corresponding table with all values in the histogram can be found in Table B1.

White Share Decline by Distance from Minority Cluster Boundary
Chicago MSA, 1970–1980, Cluster Size = 1




Figure 11: This figure shows drastic changes in the White share in percentage points by location for Chicago 1970-1980. Change is largely concentrated at the boundary of clusters, and the more drastic the change, the more concentrated it is close to the boundary. Defining tipping as drastic changes in racial shares shows that this happens at the boundary of clusters in Chicago in this period.




Figure 12: Drastic change is concentrated at the boundaries of clusters. Change in White Share of 25 p.p. or more, Chicago MSA, 1970-1980. Panel (a) is Chicago MSA, panel (b) is Chicago central city. Tracts with change of 25 p.p. or more in White share are shaded yellow.

White Share Decline by Distance from Minority Cluster Boundary
 All MSA, 1970–1980, Cluster Size = 1




Figure 13: White share decline by distance bin for all MSAs, 1970-1980. The higher the cutoff to define drastic racial change, the more is it concentrated at the boundary of clusters. The number above each bar indicates the number of tracts in that bar, with the corresponding share on the y -axis. Defining tipping as drastic changes in racial shares, this figure demonstrates that tipping for all MSAs is concentrated at the boundary of racial clusters.

When we examine this for all MSAs in this period, but vary the cutoffs, we see again that change is centered on the tracts at the boundary of the Minority cluster and that this concentration increases as the change considered becomes more drastic.

3.5 Summary of Empirical Evidence

Overall we interpret our empirical evidence as being astonishingly well in line with the predictions of our simulated spatial proximity model. Racial clusters are an ubiquitous feature of US cities across all time periods of our investigation (Hypothesis 1). We find strong support that Minority clusters are highly segregated and that White clusters used to be highly segregated but have experienced an increase in their median Minority share between 1970 and 2000 (Hypothesis 2). Rental prices and home values show clear patterns that are in line with the predictions of a model that features spatial spillovers and racial preferences (Hypothesis 3). Finally, racial change predominantly occurs at the boundary of racial clusters and tipping in the sense of drastic changes is particularly concentrated in these boundary regions (Hypotheses 4 and 5).

We would like to emphasise that the evidence we provide for the relevance of spatial spillovers

is not causal, but rather strongly suggestive. Other patterns, such as spatially correlated labor market access, different incomes and wealth, or discrimination in the housing market – all varying by race – likely explain part of the patterns observed in the census data. Nevertheless, we argue that spatial spillovers provide such a simple and sensitive explanation for what we observe that they must be key in rationalizing the cross-section and predicting where long-term racial changes are happening.

4 Empirics of the Bounded Neighborhood Model

One of the main predictions of the spatial proximity model is that drastic racial changes and tipping are concentrated at the boundaries of racial clusters. In our empirical investigation focusing on decadal changes in census tracts’ racial composition we found strong supportive evidence for this pattern.

One of the most important and widely cited contributions to the empirical literature on neighborhood tipping is Card et al. (2008). The authors develop a reduced form approach to identify tipping points at the MSA-level from census data. They find significant tipping points for many cities, compare the location of tipping points across cities and track their development over the decades from 1970 until 2000. In their section on the geography of tipping, Card et al. write that “*Taken together, these results are not consistent with the predictions of the expanding ghetto model. Tipping effects are, if anything, strongest far from the existing ghetto. We conclude that this model cannot account for the nonlinear dynamics we see [...]*” (Card et al., 2008, p. 205). The authors thus argue that tipping is not concentrated at the boundaries but that it is occurring far from existing Minority clusters (which they refer to as the ghetto).

How did Card et al. (2008) end up with a conclusion regarding the locus of racial neighborhood tipping that is so fundamentally different from ours? To understand their results better we revisit their methods and findings in this section. As we are going to show, their empirical results are driven by (1) the usage of global polynomials when identifying tipping points through regression discontinuity analyses and (2) a small number of suburban census tracts that experience dramatic population growth and *White entry* during the decades that the authors investigate. Neighborhood tipping is generally understood to be characterized through *White exit* in response to increasing Minority shares. This view of tipping as a phenomenon of White exit is also true in the partial equilibrium framework that Card et al. (2008) use to motivate their empirical strategy. We show that once one uses local linear regressions in the regression discontinuity analysis and once one excludes suburban tracts with


dramatic population growth, the candidate tipping points identified in Card et al. (2008) lose their significance and become precisely estimated zeros.

4.1 A Re-examination of Chicago

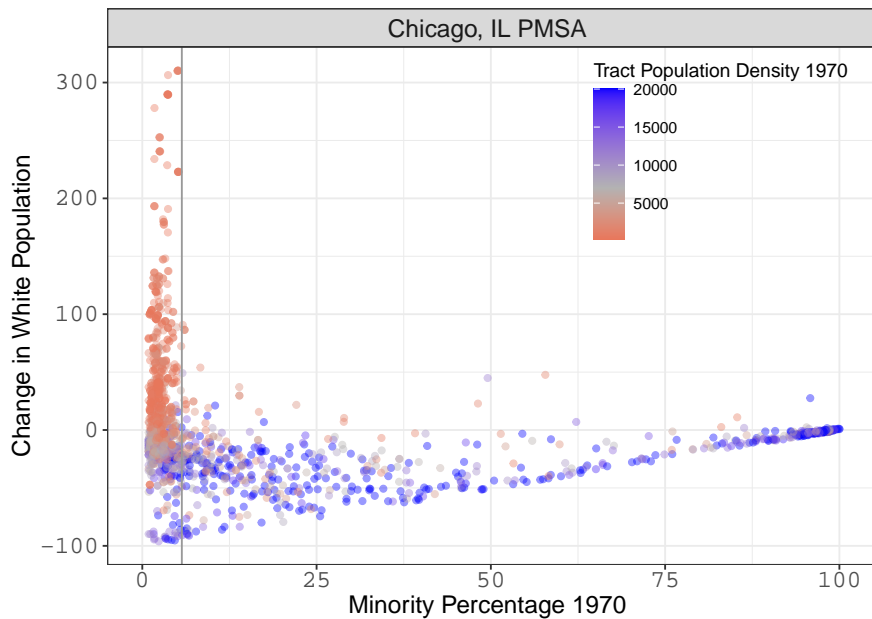
We start our investigation of Card et al. (2008) results with a re-examination of Chicago and replicate their striking Figure 1 which is our Figure 14a. It plots the binned change in the White population 1970-1980 as a share of the tract’s 1970 total population on the vertical axis against a tract’s initial Minority share in 1970 on the horizontal axis.²² Card et al. provide two different methodologies to identify candidate tipping points on the MSA-level. The vertical line in figure 14a is the posited tipping point for Chicago using their preferred fixed point method. It is 5.7% in 1970. The horizontal line is the unconditional mean for the change in White population. Allowing these lines to define quadrants, we see that the first and third quadrants are entirely empty. The jump at the tipping point s_{mj}^* is measured as the discontinuity in a quartic polynomial local linear regression in the deviation of the Minority share m from the posited tipping point, implemented separately to the left and right. For Chicago in the 1970-1980 period, this jump at the tipping point exceeds 30%.

Naturally, the binning shrouds heterogeneity. But the heterogeneity may provide insight to the forces identifying the discontinuity at the tipping point. So we next examine the unbinned data for Chicago, which we show in Figure 14b. Several features emerge clearly. Looking to the right of the tipping point s_m^* , there are now only a quite modest number of observations in the first quadrant, suggesting that those observations that featured growth in the White population had a likewise quite modest influence on the binned means in that region. The contrast for points below s_m^* leaps out. The third quadrant, wholly empty in the binned data, is actually quite heavily populated with observations with great losses in White population. The fact that for initial Minority share s_m below s_m^* , the binned means of White share growth are all positive implies that the presence of this large number of tracts in the third quadrant with declining White shares is outmatched by spectacular growth of the White share in tracts in the second quadrant. Part of this particular result, though, is mechanical. The White population relative to the initial tract population is bounded below by 100%, while there is no upper bound for positive values of White percentage growth (which reaches above 300% for some tracts).

²²To be precise, the y -variable is defined as “(tract-level White population in 1980 minus tract-level White population in 1970) divided by tract-level total population in 1970”. The units for this variable are thus percentage point change in White population relative to 1970; tract geographies are standardized at 2000 Census geographies to allow for cross-decadal comparisons. This y -variable is the same as used in Card et al. (2008); we maintained their restrictions to the set of tracts in the estimation sample, which are explained in-depth in their paper.



(a)



(b)

Figure 14: (a) Binned Data for Chicago and (b) Unbinned Data for Chicago, 1970-1980. Subplot (a) contains 100 scattered points of width 1 percent (in terms of Minority percentage) which represent the average change in White population for all tracts in that Minority percentage band. A kernel mean smoother is overlaid. Subplot (b) illustrates the unbinned, raw data. Tract population densities were top-coded at 20,000 and bottom-coded at 200, with an overall median of around 7,000. The vertical line represents the fixed-point estimated tipping point for Chicago in 1970 (5.7%). Subplot (a) figure replicates Figure 1 from Card et al. (2008), using their data. We see, in (b), that among tracts below the posited tipping point, tracts whose White population grew were almost exclusively initially low population density areas.

From a theoretical perspective, these features of the unbinned data are worrying. The tipping model of a single tract developed in Card et al. (2008) is about drastic *exit* of Whites from a tract. If you contrast it in the same setting with a single tract that does not tip, this should be about a contrast between two tracts, *both* of which experience White exit. That is, the contrast is whether this White exit is drastic or marginal. The model has no story about tracts with growth in White population. There is no general equilibrium in the simple bounded neighborhood model. One can ponder where departing Whites may go, but this strictly has no role at all in tipping point theory.

These features suggest that patterns of growth in White tracts may play a critical role in identification in the canonical empirics of tipping, and motivates a closer look at this aspect. One natural dimension to explore, especially given the mechanical aspects of the growth in White share in the second quadrant, is the role of the base 1970 tract population density in giving rise to the contrasts above and below the posited tipping point. We examine this by coloring tracts based on their population density in 1970²³. In the base year of 1970, blue tracts are high population density tracts and red tracts are low population density tracts.

Focusing attention on tracts below the posited tipping point, hence in the second and third quadrants, there appears to be a strong tendency for those losing White population to initially be high population density tracts, and those gaining White population to initially be low population density tracts. When we turn to regression analysis, this will motivate a couple of amendments to standard approaches to estimating the tipping points. One is to introduce population weights, to be sure that all of the results are not being driven by explosive growth in very sparsely populated locations. The second, to focus attention on the direct predictions of tipping theory, is to explore the consequences of restricting the sample to theory-consistent tracts, which do not have such explosive growth.

Our results suggest that explosive growth of the White population in initially low population tracts may play an important role in the standard results claiming to identify tipping. However, this suggests making explicit two contrasting hypotheses about what may be giving rise to these results.

- The tipping hypothesis predicts that in comparing tracts with initial Minority share $m_0 < m^*$ and those with initial share $m_0 > m^*$, both experience White exit, but in the former case it is limited, as they do not tip, while in the latter case it is drastic.

²³Population density is a more robust statistic to use in this situation than population levels, as we have fixed census tract geographies to 2000-vintage boundaries.

Location *per se* of the tracts should play no role, as this has no role in the tipping (bounded neighborhood) theory.

- We will refer to the competing hypothesis as “biased White suburbanization.” One critical element is that this theory allows for the possibility of tipping, but does not require it. Even if White departure from many central city tracts was due to shocks *in a non-tipping world*, those Whites may nonetheless seek out new locations in initially low density and low initial Minority share suburban neighborhoods. Location, under this hypothesis, is crucial for understanding the evolution. But the presence of many of these tracts in the category of $m_0 < m^*$ is not about the type of tracts that Whites exit, but instead about the kind of tracts they enter. It is about White population growth, not decline.

While we have described the locational dimensions of the problem, we can relate it more directly to identification in the standard tipping approach. Recall that Figure 14b divided the tracts into four quadrants, where quadrants 2 and 3 are below the tipping point and quadrants 1 and 4 are above it. We can produce an analogous figure using maps to illustrate geographically which tracts appear in which quadrant. The resulting maps appear as Figure 17.

The tipping hypothesis would suggest focusing on the contrasting outcomes between tracts that lose Whites, but are below or above the tipping point – the contrast between quadrants three and four. There are differences, but nothing that strongly suggests that those in quadrant four (i.e. above the tipping point) had particularly strong declines in White population relative to those in quadrant three.

By contrast, the biased White suburbanization hypothesis would focus on the contrast between locations with White growth that were either below the posited tipping point or above it – the contrast between quadrants two and one. The maps demonstrate that there is nothing subtle at all. Quadrant one has almost no mass. That is, tracts that gained Whites were rarely those with initial Minority share above the posited tipping point. Overwhelmingly, if a tract in Chicago 1970-1980 gained White population, then it was a low population density, low initial Minority share tract.

The results illustrated in Figure 17 thus provide powerful *prima facie* evidence in support of the biased White suburbanization hypothesis in place of the tipping point hypothesis. Of course, this does not tell us that tipping points do not exist. It only suggests that the existing analysis may be claiming that results support the tipping point hypothesis when they

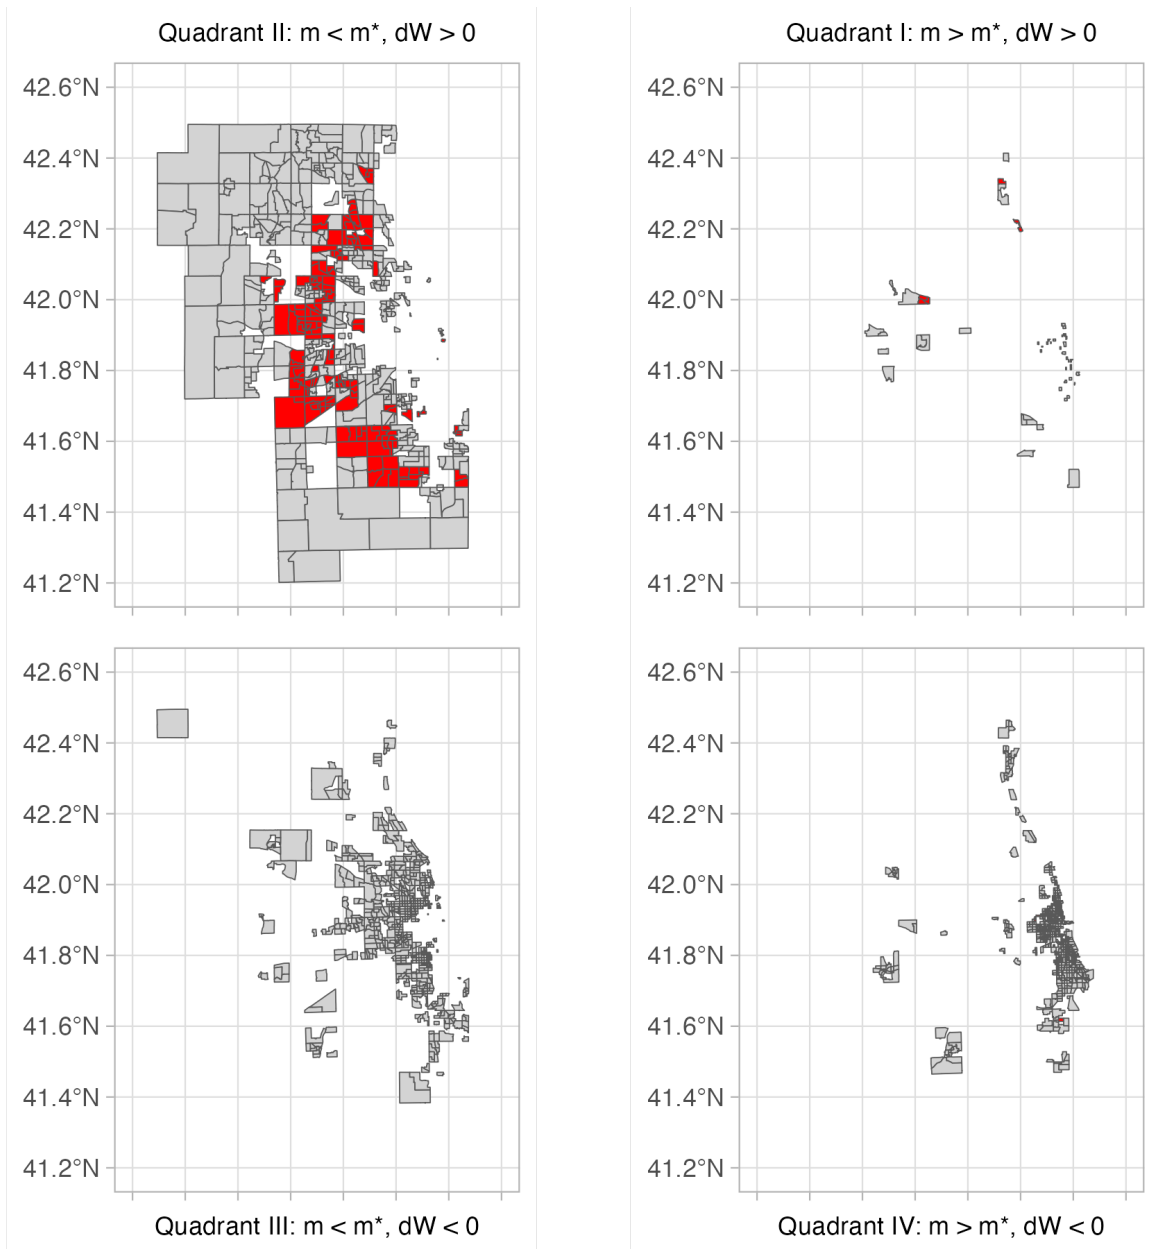


Figure 15: Mapping the tipping point: driven by biased White suburbanization. Here we take the four quadrants defined by the unbinned data for Chicago in Figure 14b and plot the tracts for each quadrant on its own map. Red shading means the tract experienced 60 percent or larger population growth across the decade. These maps by quadrant highlight the critical role in quadrant 2 of high growth suburban tracts in driving what in Card et al. (2008) appeared to be tipping, but is instead biased White suburbanization.

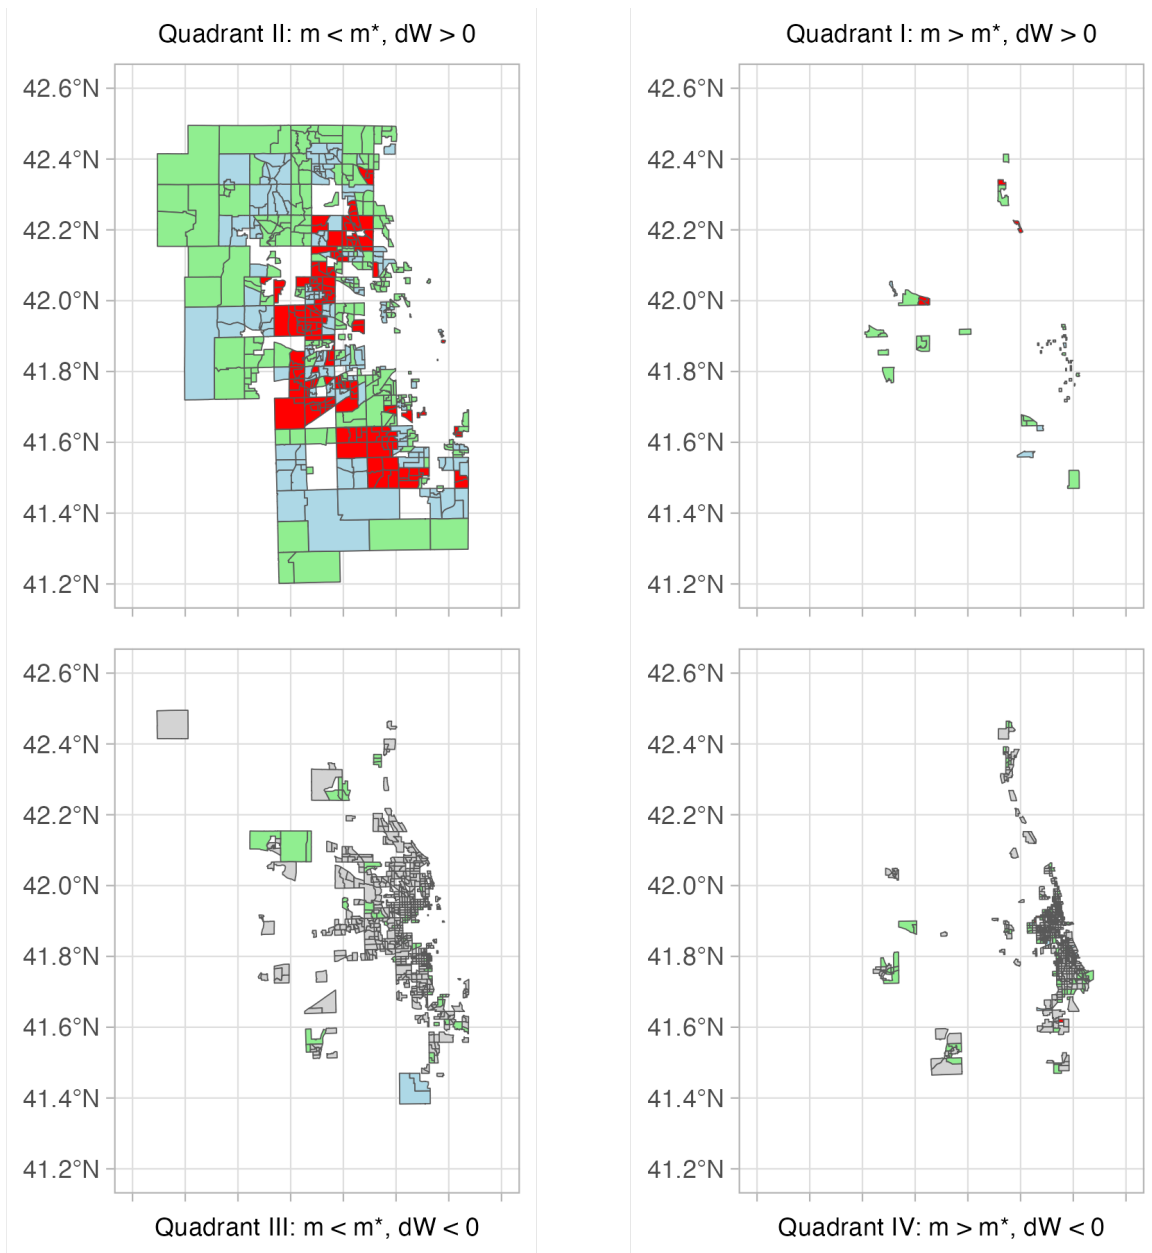


Figure 16: Mapping the tipping point: driven by biased White suburbanization. Here we take the four quadrants defined by the unbinned data for Chicago in Figure 14b and plot the tracts for each quadrant on its own map. Red shading means the tract experienced 60 percent or larger population growth across the decade; light blue is between 30 and 60 percent; light green is less than 30 percent. These maps by quadrant highlight the critical role in quadrant 2 of high growth suburban tracts in driving what in Card et al. (2008) appeared to be tipping, but is instead biased White suburbanization.




Figure 17: Mapping the tipping point: driven by biased White suburbanization. Here we take the four quadrants defined by the unbinned data for Chicago in Figure 14b and plot the tracts for each quadrant on its own map. Shades of red indicate low population density in 1970, while shades of blue indicate high population density in 1970. These maps by quadrant highlight the critical role in quadrant 2 of high growth suburban tracts in driving what in Card et al. (2008) appeared to be tipping, but is instead biased White suburbanization.

appear instead to be driven by a distinct hypothesis, that of biased White suburbanization. We should also examine these results in a regression analysis, where it will be straightforward to test the competing hypotheses using our more extensive sample of cities.

4.2 Tipping and Biased White Suburbanization: Regressions Across MSAs

The Chicago example has raised questions about both the robustness and interpretation of the Card et al. (2008) results claiming to find strong evidence of racial tipping. Here, we will explore these concerns in a few steps. The first is to note that preferred regression discontinuity designs have evolved since their paper, and so we need to implement a modern approach. The second takes note of the competing hypothesis that the measured “tipping” is actually driven by White suburbanization, and considers two approaches to adjusting for this. The first approach simply adds population weighting to reduce the influence of tracts that are initially very low population, but which partly as a result may experience spectacular growth rates. The second approach is what may be termed a “theory-consistent” approach. This notes that the theory of tipping, partial equilibrium and focused on a single tract, has no explicit prediction about tracts with White population growth. It thus considers restricting the sample to tracts where the total population growth is less than a fixed threshold.

We begin by replicating the original results in Card et al. (2008). They have two methods for identifying candidate tipping points, one based on a model with a single structural break, and their preferred fixed point approach, the Minority share at which the tract population grows at the city-wide average. Having used a fraction of their data to identify candidate tipping points, the remainder of the data is used to estimate the magnitude of the jump at the discontinuity using a regression discontinuity design with a quartic polynomial in deviations from the candidate tipping point. In their preferred fixed point approach, illustrated in the first column of Figure 18, the jumps in the decades respectively of 1970, 1980, and 1990 are all significant at -12 , -14 , and -7 percent. They get broadly similar, although somewhat less precise, results with their structural break method.

The RD design that Card et al. (2008) implemented, relying on a global quartic polynomial, is now in disfavor.²⁴ Thus it is crucial to check that their results survive introduction of

²⁴Gelman and Imbens (2019) have the helpful title “Why high-order polynomials should not be used in regression discontinuity designs.” They state (p. 447) “It is common in regression discontinuity analysis to control for third, fourth, or higher-degree polynomials of the forcing variable. There appears to be a perception that such methods are theoretically justified, even though they can lead to evidently nonsensical results. We argue that controlling for global high-order polynomials in regression discontinuity analysis is a flawed approach with three major problems: it leads to noisy estimates, sensitivity to the degree of the polynomial, and poor coverage of confidence intervals. We recommend researchers instead use estimators based on local linear or quadratic polynomials or other smooth functions.”




Figure 18: Summary figure of regressions using Card et al. (2008) data to estimate tipping point discontinuities across decades and estimation methods. Global regressions use a quartic polynomial, as in Card et al. (2008), while local linear regressions rely on optimal bandwidth selection following Calonico et al. (2020). Confidence bands are given for the 95% level. In the last two columns, text in parentheses details what fraction of the sample population is removed by the respective restriction. The results from Card et al. (2008) about the negative, significant impact on white tract population beyond the tipping point go to insignificance or precise zeros.

more modern methods. In the following, we use a local linear regression with a triangular kernel including the same controls and MSA-level fixed effects as Card et al. (2008). The kernel bandwidths are chosen following the procedures detailed in Calonico et al. (2014, 2019, 2020).²⁵ We show the resulting estimates for the fixed point method, but with local linear regressions, in Figure 18. The point estimate for the decade from 1970 remains unchanged at -12 percent. But the results for 1980 and 1990 have moved toward zero and are insignificant.²⁶ In short, the introduction of modern regression discontinuity methods alone has importantly weakened the results, leading to negative and significant results only in 1970 and not in the other decades.

We now pair the modern RD method with population weights, as a first step to reducing the influence of initially low population tracts, which as a result may enjoy spectacular growth rates. This appears in the fourth column of Figure 18. For 1970, relative to the unweighted local linear regression, the point estimate falls by half and becomes insignificant, although not precisely measured. For 1980 and 1990 the coefficients remain insignificant and close to zero.

Finally, we pair the modern RD method with a theory-consistent restriction on the growth rate of tracts included in the sample for varying thresholds. Since tipping theory has no prediction about tracts that will experience growth in the White population, a sample selection that caps the admissible values for total population growth should not affect whether we observe tipping in the remaining tracts.²⁷ The results are reported in the right side of Figure 18 for three decades and two different population growth caps. In 1980 and 1990, this experiment leaves the values as statistical zeros. We see that in 1970, the one remaining significant negative jump of -12 percent becomes a precisely measured zero when we restrict attention to tracts whose growth in the decade is 60 percent or less. For this latter period MSA level average population growth was only 13% and only 7% of tract-level observations exceeded the 60% cap, leading us to conclude that this is not an overly stringent restriction on the sample.

²⁵We use the R package `rdrobust` and report results from the bias-corrected effect estimate (for details see Calonico et al., 2015).

²⁶The results with the local linear regressions and their structural break method are broadly in parallel. There is a significant (and even larger) negative jump of -22 percent in 1970, but smaller and insignificant jumps in 1980 and 1990.

²⁷While we have emphasized that the restriction is consistent with the aspatial common threshold theory that *all* neighborhoods should be losing White population, one can also think about this in the context of the key result of Boustan (2010), which finds causal estimates of 2.7 White departures for each Black arrival. If this is viewed as tipping, it is hard to see how this would be consistent with the phenomenon relying on high levels of *entry* to tracts.

4.3 Has Tipping Gone Kerplunk?!

Has tipping gone kerplunk? In a word, no. The work of Card et al. (2008) is a brilliant implementation of a variant of the Schelling tipping model. Developing an approach that allows the identification of tipping points without explicitly specifying a structural model and then causally estimating racial preferences is extremely appealing. Part of the beauty of the paper is the apparently robust results in spite of the simplicity of the model and the highly stringent maintained assumptions required to take it to data. But the results proved partially vulnerable to a fate that haunts every researcher: methods evolve. The introduction of modern regression discontinuity methods alone eliminated the significance of their tipping results for all MSAs for 1980 and 1990, leaving only 1970 as a period of significant negative effects of tipping.

Our example of Chicago for 1970 indicated another vulnerability of working with their simple reduced form strategy motivated by a partial equilibrium model of a single neighborhood. When general equilibrium considerations are introduced, so adding-up constraints bind, there is an alternative interpretation of what is driving their results. In this case, we have shown that the apparent result showing drastic White *exit*, i.e. their approach to tipping, was actually largely driven by the analytically distinct biased White suburbanization, as Whites *entered* suburban tracts that had both low initial population densities and low Minority shares, leading to explosive White population growth rates in the latter locations.

While the precise Card et al. (2008) approach to tipping proved not to be robust, this does not imply that tipping of neighborhoods is not a real and important phenomenon. To estimate tipping points we would however caution against the reduced form approach. Instead, we think that estimating racial preferences causally and using a calibrated structural model to identify tipping points that can vary across locations is the right approach. We argue that a component of such a model should be racial preference spillovers across neighborhoods as they are key to rationalize observed patterns of segregation and make sense of where drastic racial change is likely to happen. Given that drastic racial evolution in our framework is consistent with both the existence and absence of formal bifurcations, we prefer to define tipping directly as the drastic change in racial shares.

5 Conclusion

A recent empirical literature identifies important causal links between neighborhood characteristics and economic outcomes through both experiments (Moving to Opportunity) and quasi-experiments (Atlas of Opportunity). This provides strong motivation for moving be-

yond pure observational studies or measure-zero experiments to consider interventions at scale that might provide greater opportunities to those currently lacking such opportunities. If such policies are to be implemented at scale, it will be crucial, then, to understand this from a general equilibrium perspective. It is all the more important when the interventions will also change the racial composition of neighborhoods, as the literature on racial change suggests the possibility of racial tipping that may offset or undo the intended effects of policy.

In spite of the vast existing social science literature on residential neighborhood segregation, we still have surprisingly little empirical research in economics that helps us to understand the cross-sectional and dynamic response of neighborhoods to spatial racial shocks. We advance understanding in these dimensions.

We begin with a discrete choice model with tastes for race, but that in contrast to most of the prior literature allows these tastes to spill past the individual location to neighboring locations. We use this model for simulations that allow us to generate hypotheses that we can examine for US metropolitan areas. We develop six such hypotheses that touch on the importance of racial clusters in US metro areas; the patterns of segregation interior to and at the edge of these clusters; housing price gradients inside and at the edge of clusters; the location of racial change at the boundaries of clusters; the focus of drastic change, or tipping as we define it, also at the boundary of clusters; and that in the face of racial shocks, Minority clusters expand near to existing clusters and Whites move to remote, low density, low Minority share neighborhoods in the suburbs. An examination of our data for US metros between 1970-2000 provides strong support for these hypotheses.

Our emphasis on the locus of drastic change at the boundary of existing Minority clusters directly conflicts with core results in the influential work of Card et al. (2008). In order to reconcile our results with theirs, we re-examine their results through the lens of our work. We illustrate this first for a core case that they examine, that of Chicago in the period 1970-1980. They developed results that they interpret as a tipping point at a Minority share of 5.7% within a tract, revealing a jump of minus 30 percentage points for tracts beyond this tipping point. We demonstrate that this is not a convincing interpretation of their own data. The tipping model posits that the differential growth in the White population should be driven by White exit from tracts that exceed the tipping threshold. Instead, in the case of Chicago in this period, the measured tipping is driven by White entry to initially low density, low Minority suburban tracts. There is no evident discontinuity at the tipping point among tracts that lose White population.

We also examine this in a regression framework that builds on Card et al. (2008). We introduce two new elements. One is that we update their methods for regression discontinuity. We also introduce theory-consistent caps on the extent of total population growth allowed for inclusion in the sample. Just introducing modern methods leads the results in the latter two decades of their sample to become insignificant. Consistent with theory, capping allowable total population growth at 60%, and combining this with modern methods leads their results on tipping for all MSAs in the periods 1970-2000 instead to become precisely measured zeros.

Our results lead us to believe it is fruitful to change how we define tipping. The focus on formal bifurcations at an MSA-level common tipping point of Card et al. (2008) proved fragile. Our approach demonstrates that drastic local racial change is possible both when formal bifurcations exist and when they do not, since the presence of sufficiently strong homophily can give rise to elastic responses in either case. For this reason, we prefer to define tipping directly as drastic racial change. Consistent with the spatial proximity model of Schelling (1971), tipping under this definition is concentrated at the boundaries of racial clusters.

In sum, we have developed a framework based on a discrete choice model with spatial racial spillovers that identifies a set of hypotheses strongly consistent with the data for the US in the period 1970-2000. We believe it provides insight to understanding both the cross-sectional and dynamic features of racial segregation.

References

- Almagro, M., Chyn, E., and Stuart, B. A. (2023). Urban Renewal and Inequality: Evidence from Chicago’s Public Housing Demolitions.
- Axtell, R. L. and Farmer, J. D. (2022). Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature*.
- Bagagli, S. (2023). The (Express)Way to Segregation: Evidence from Chicago.
- Bayer, P., Casey, M. D., McCartney, W. B., Orellana-Li, J., and Zhang, C. S. (2022). Distinguishing causes of neighborhood racial change: A nearest neighbor design. Technical report, National Bureau of Economic Research.
- Bayer, P., Ferreira, F., and McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of political economy*, 115(4):588–638.

- Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2016). A Dynamic Model of Demand for Houses and Neighborhoods. *Econometrica*, 84(3):893–942.
- Bayer, P. and Timmins, C. (2005). On the equilibrium properties of locational sorting models. *Journal of Urban Economics*, 57(3):462–477.
- Bayer, P. and Timmins, C. (2007). Estimating Equilibrium Models of Sorting Across Locations. *The Economic Journal*, 117(518):353–374.
- Becker, G. S. and Murphy, K. M. (2000). *Social economics: market behavior in a social environment*. Belknap Press of Harvard University Press, Cambridge, Mass.
- Blair, P. (2023). Beyond Racial Attitudes: The Role of Outside Options in the Dynamics of White Flight. NBER working paper w31136, National Bureau of Economic Research, Cambridge, MA.
- Boustan, L. (2010). Was postwar suburbanization 'white flight'? evidence from the black migration. *Quarterly Journal of Economics*.
- Caetano, G. and Maheshri, V. (2017). School segregation and the identification of tipping behavior. *Journal of Public Economics*, 148:115–135.
- Caetano, G. and Maheshri, V. (2023). A Unified Empirical Framework to Study Segregation.
- Calonico, S., Cattaneo, D., M., and Titiunik, R. (2015). rdrobust: An R Package for Robust Nonparametric Inference in Regression-Discontinuity Designs. *The R Journal*, 7(1):38.
- Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2020). Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs. *The Econometrics Journal*, 23(2):192–210.
- Calonico, S., Cattaneo, M. D., Farrell, M. H., and Titiunik, R. (2019). Regression Discontinuity Designs Using Covariates. *The Review of Economics and Statistics*, 101(3):442–451.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals. *Econometrica*, 82(6):2295–2326.
- Card, D., Mas, A., and Rothstein, J. (2008). Tipping and the dynamics of segregation. *Quarterly Journal of Economics*, 123(1).
- Chetty, R. and Hendren, N. (2018a). The Impacts of Neighborhoods on Intergenerational Mo-

- bility I: Childhood Exposure Effects*. *The Quarterly Journal of Economics*, 133(3):1107–1162.
- Chetty, R. and Hendren, N. (2018b). The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates*. *The Quarterly Journal of Economics*, 133(3):1163–1228.
- Chetty, R., Hendren, N., and Katz, L. F. (2016). The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review*, 106(4):855–902.
- Chyn, E. (2018). Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children. *American Economic Review*, 108(10):3028–3056.
- Chyn, E., Haggag, K., and Stuart, B. A. (2022). The Effects of Racial Segregation on Intergenerational Mobility: Evidence from Historical Railroad Placement.
- Chyn, E. and Katz, L. F. (2021). Neighborhoods Matter: Assessing the Evidence for Place Effects.
- Couture, V., Gaubert, C., Handbury, J., and Hurst, E. (2023). Income Growth and the Distributional Effects of Urban Spatial Sorting. *The Review of Economic Studies*, page rdad048.
- Cutler, D. M., Glaeser, E. L., and Vigdor, J. L. (1999). The rise and decline of the american ghetto. *Journal of political economy*, 107(3):455–506.
- Dai, T. and Schiff, N. (2023). The structure and growth of ethnic neighborhoods. *Journal of Urban Economics*, 137:103570.
- Davis, M. A., Gregory, J., and Hartley, D. A. (2023). Preferences over the Racial Composition of Neighborhoods: Estimates and Implications.
- Davis, M. A., Gregory, J., Hartley, D. A., and Tan, K. T. K. (2021). Neighborhood effects and housing vouchers. *Quantitative Economics*, 12(4):1307–1346. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/QE1664>.
- Easterly, W. (2009). Empirics of strategic interdependence: the case of the racial tipping point. *The BE Journal of Macroeconomics*, 9(1).
- Ferreira, F. V., Kenney, J. H., and Smith, B. (2023). Household Mobility, Networks, and Gentrification of Minority Neighborhoods in the US.

- Gelman, A. and Imbens, G. (2019). Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, 37(3):447–456.
- Glaeser, E. and Vigdor, J. (2012). The end of the segregated century. *Manhattan Institute for Policy Research, January*, pages 23–26.
- Massey, D. S. and Denton, N. A. (1988). The dimensions of residential segregation. *Social forces*, 67(2):281–315.
- O’Sullivan, A. (2009). Schelling’s model revisited: Residential sorting with competitive bidding for land. *Regional Science and Urban Economics*, 39(4):397–408.
- Schelling, T. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*.
- Schelling, T. C. (1969). Models of segregation. *The American economic review*, 59(2):488–493.
- Tsivanidis, N. (2023). Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá’s TransMilenio.
- Weiwu, L. (2023). Unequal Access: Racial Segregation and the Distributional Impacts of Interstate Highways in Cities.
- Zhang, J. (2004). Residential segregation in an all-integrationist world. *Journal of Economic Behavior & Organization*, 54(4):533–550.
- Zhang, J. (2011). Tipping and residential segregation: a unified schelling model. *Journal of Regional Science*, 51(1):167–193.

A Additional Figures




Figure A1: Construction of bid-rent curves from demand functions. Green curves show demand for tract j at different hypothetical Minority shares s_{rj} . Red curve shows the resulting bid-rent function.




Figure A2: Weights for a Neighborhood in the Bounded Neighborhood and Spatial Proximity Models




Figure A3: The Location of tracts relative to the boundary with the Minority cluster in Chicago MSA, 1970. Panel (a) is the full MSA, and Panel (b) zooms in on the city of Chicago.

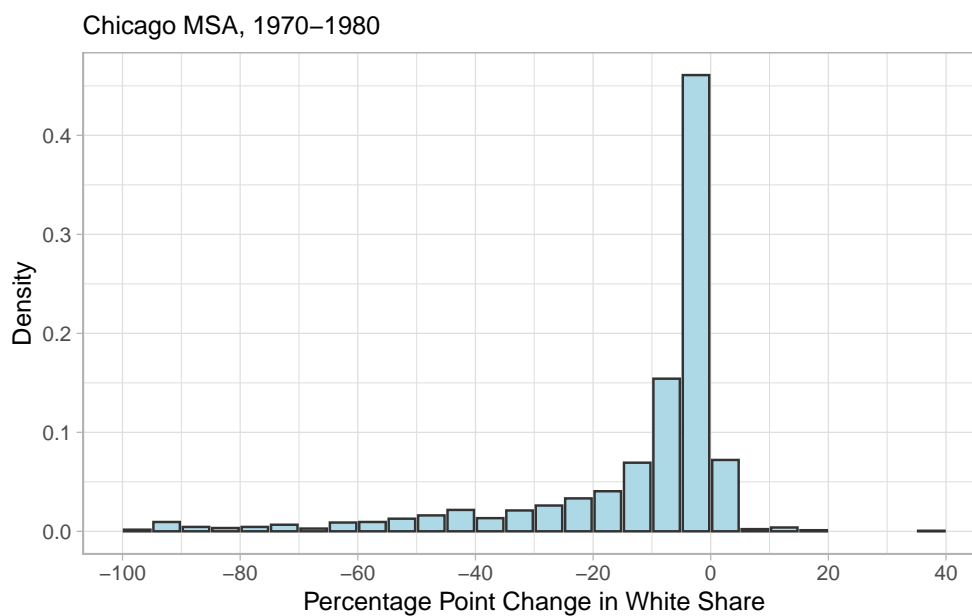


Figure A4: Density of Percentage Point Change in White Share in Chicago MSA, 1970-1980.

Relative Property Values by Distance from Minority Cluster Boundary
Chicago MSA

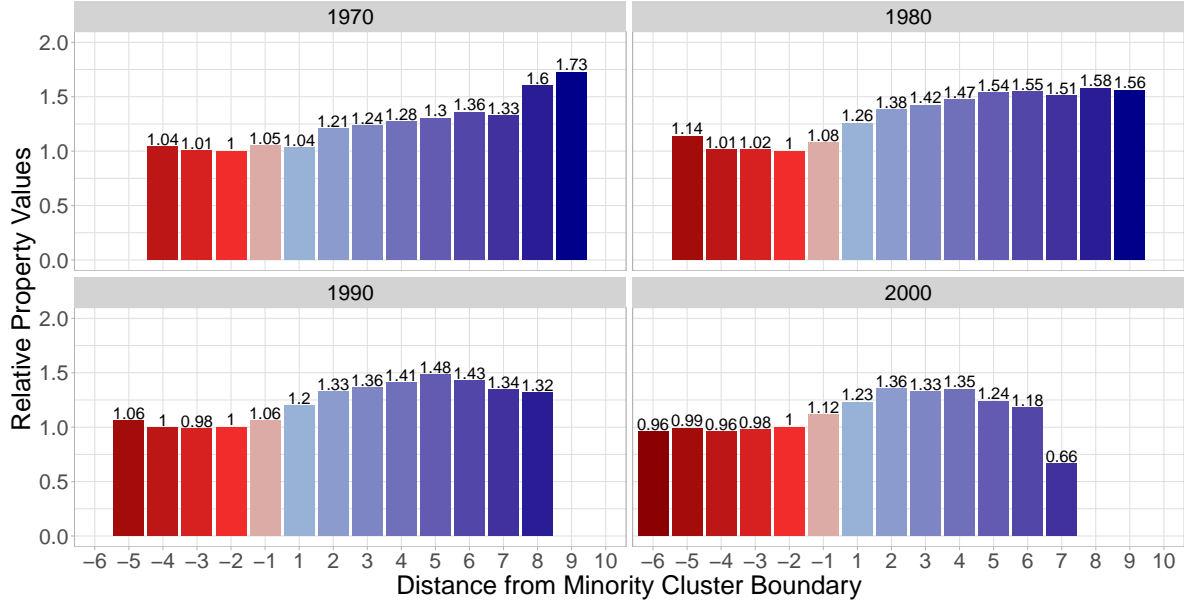


Figure A5: Median House Values relative to Values in Tracts with a Distance from the Minority Cluster Boundary of -2, Chicago MSA.

B Additional Tables

Table B1: This table shows the share of all tracts in Chicago in 1970-1980 at each percentage point change in their White share.

Density of Percentage Point Change in White Share	
Chicago MSA, 1970-1980	
Change (p.p.)	Share of Tracts (%)
(-100,-95]	0.2
(-95,-90]	0.9
(-90,-85]	0.4
(-85,-80]	0.3
(-80,-75]	0.4
(-75,-70]	0.7
(-70,-65]	0.3
(-65,-60]	0.9
(-60,-55]	0.9
(-55,-50]	1.3
(-50,-45]	1.6
(-45,-40]	2.2
(-40,-35]	1.3
(-35,-30]	2.1
(-30,-25]	2.6
(-25,-20]	3.3
(-20,-15]	4.0
(-15,-10]	6.9
(-10,-5]	15.4
(-5,0]	46.1
(0,5]	7.2
(5,10]	0.2
(10,15]	0.4
(15,20]	0.1