

Segregation That No One Seeks

Ryan Muldoon, Tony Smith, Michael Weisberg

December 3, 2010

Abstract

This paper examines a series of Schelling-like models of residential segregation, in which agents prefer to be in the minority. We demonstrate that as long as agents care about the characteristics of their wider community, they tend to end up in a segregated state. We then investigate the process that causes this, and conclude that the result hinges on the similarity of informational states amongst agents of the same type. This is quite different from Schelling-like behavior, and suggests (in his terms) that segregation is an instance of macro behavior which can arise from a wide variety of micro motives.

Weak, individually-held preferences can be significantly amplified when aggregated in a population. This fact is exemplified in simple agent-based models of segregation, that show how weak preferences to have like individuals as neighbors leads to far greater segregation than any individual desires.

In his landmark studies of this phenomenon¹, Thomas Schelling represented a neighborhood as a grid with two types of agents placed randomly on it. Each agent was allowed to move to nearby unoccupied cells in order to satisfy its preference to be like at least thirty percent of its neighbors. Despite the relatively low preference for having neighbors of the same type, these models exhibit significant segregation and clustering of the agents. This phenomenon has become known as the *Schelling result* or *Schelling segregation*.

¹Schelling (1971)

Subsequent theoretical work has investigated a wide range of neighborhood definitions and agent utility functions. In every case that we are aware of, researchers have found Schelling-like models to exhibit segregation. Since the segregation result is extremely robust against changes to utility function and neighborhood definition, we investigated whether segregation could survive an even more strenuous perturbation: agents who explicitly aim to be in the smallest minority. We will show that when model-individuals know about their communities in addition to their neighborhood and strictly prefer to be in the minority in their communities, widespread segregation can develop.

Our investigation employs a series of Schelling-like models where individuals placed randomly on a toroidal grid assess their utility on the basis of how many similar individuals are in a viewable radius. The utility functions and radii vary in the models, with the result that even agents who prefer to be in the smallest minority end up highly segregated when their viewable radii are of intermediate size relative the grid. These results demonstrate the extreme fragility of integrated populations because they show that even when individuals try as hard as possible to be in the minority, imperfect information about their environment actually leads them to segregate.

1 Schelling’s Models of Segregation

Thomas Schelling famously asked “what leads a neighborhood to segregate?” He showed that racial segregation can occur even when no explicit racism is present by constructing a physical agent-based model of a population. In this model, dimes and nickels represented two types of individuals, A and B , and the squares on a chess board represented spatial location. Each individual prefers that 30% of its neighbors be of the same type. So the A s want 30% of their neighbors to be A s and likewise for the B s. Schelling’s neighborhoods were defined as standard Moore neighborhoods, a set of nine adjacent grid elements. An agent standing on some grid element e can have anywhere from zero to eight neighbors in the adjoining elements.

Although Schelling didn’t explicitly provide a utility function, the pref-

erence described above is usually interpreted to mean that each agent is indifferent between having 30-100% of her neighbors be alike, but finds having fewer than 30% of her neighbors be alike unacceptable. Because of the constraints of a grid's geometry, in the case of a full neighborhood, the preference boils down to wanting to have at least three of one's eight neighbors be alike and to equally prefer 3-8 like neighbors.

The dynamics of Schelling's model involve agents sequentially choosing to remain in place or move to a new location. When it is an agent's turn to make a decision, it determines whether there is a sufficient ratio of alike agents amongst its neighbors. If this condition is met, the agent is satisfied and remains where it is. If it is not, the agent then moves to the nearest empty location. This sequence of decisions continues until all of the agents are happy where they are, and do not try to move.

What one notices from watching this model unfold is that there is a contagion effect: agents that were originally satisfied can become disgruntled as soon as a neighbor leaves or a new one moves in. It is in this way that the decision of a single agent can dramatically affect the decisions made by the entire population. A small patch of dissatisfaction can result in widespread movement, and ultimately, segregation of the "city." Although there are a few possible grid configurations which are fully integrated and in which every agent is happy, these are rare, and nearly impossible to arise from agent movement. The dominant equilibrium state of the model is segregation.

Schelling concluded that small preferences for similarity can lead to massive segregation. This conclusion is quite robust across many changes to the model including different utility functions², different rules for updating³, differing neighborhood sizes, and different spatial configurations⁴. These studies show that it is extremely hard to avoid segregation when agents have some preference for like neighbors. However, earlier studies have not investigated what happens when agents are committed to living in a fully integrated neighborhood. What happens when agents have maximally hetero-

²Bruch and Mare (2006); Panes and Vriend (2007); Zhang (2004)

³Bruch and Mare (2006)

⁴Fossett and Dietrich (2009)

geneous preferences? What happens when they want the maximum possible diversity among their neighbors?

This paper takes up these questions. We will investigate them by introducing a new class of Schelling-like models. But before turning to these models, let us consider how the situation might play out in a real population. Consider a neighborhood of immigrants where the primary characteristic is spoken language. All the agents begin by determining how many Spanish, Mandarin, Japanese, Lithuanian, and English speakers are in the neighborhood. They then determine if their own language is the smallest minority, part of a tie for the smallest minority, or exactly tied with all of the others. If in a given period of time, any of these conditions are met for an individual, then that individual is happy and remains in place. If not, the individual tries to rectify the situation immediately by moving to a new neighborhood. We will try to capture this situation in a set of models called *maximally heterogeneous preference models* (MHP models).

The general description of MHP models given above leaves considerable latitude about how neighborhoods are defined. While it is possible to use the nine-neighbor Moore neighborhood like Schelling, we will allow for varying neighborhood sizes because in almost all real cases, individuals interact with more than just their immediate neighbors. To take just one example, cities are often broken down into several communities, often with a number of different ethnic enclaves, and can gain reputations of having their own particular character that may be distinct from the city at large.

By using a more inclusive notion of community, we can also more realistically capture the factors individuals use in deciding to choose where to live. Individuals looking for housing do not typically consider single city blocks in isolation of everything else when looking for housing. They care about the larger area, whether it is safe, and what its amenities might be. It is very often the case that individuals first pick which community they'd like to live in, and then go about looking for suitable neighborhoods within that community.

2 MHP Models

Our investigation of the MHP utility function was conducted with a spatially explicit agent based model. This approach is very close to Schelling’s original model and will allow for direct comparison. The model begins with a grid corresponding to the largest spatial area under investigation, which we will call a *virtual city*. This virtual city is composed of cells which we can think of as addresses. In our models, there are 1,225 cells arranged on a 35 x 35 grid. This grid is wrapped around a torus to prevent edge effects. On to this grid, we randomly place 498 agents. In cases where two agents occupied the same positions, we reoriented one of the overlapping agents, by having it find a new location at random.

Each agent has some fixed property p which, following Schelling, we can think of as race, first language, or some other (unchanging) observable property. The simulations reported here all consider cases where agents have one of three possible values of this property, designated as *types* 1, 2, and 3. In all simulations, these types are evenly distributed among the 498 agents, yielding 166 agents of each type.

In addition to their intrinsic properties, each agent also has a *radius of vision*, r , within which other agents can be seen. The value of r defines the agent’s perceived *community*, which may range in size from the agent’s closest neighbors ($r = 1$) all the way up to the entire grid ($r = 17\sqrt{2}$ which is approximately 24). To keep the analysis simple, every agent is given the same value for r in each simulation. This corresponds to 498 overlapping communities, each with radius r .

Given that this is an agent-based model, each agent implements a movement strategy based on an individual utility function. If N_p^a denotes the number of agents of type p in agent a ’s community and if (i, j, k) denotes some permutation of $(1, 2, 3)$, then the *MHP utility function*, U_i^a for each

agent a of type i is defined to be a *satisficing utility*⁵ of the following form:

$$U_i^a = \begin{cases} 1 & \text{if } N_i^a < (N_i^a + N_j^a + N_k^a)/3 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Hence an agent a of type i will move to new position whenever $U_i^a = 0$.

The basic structure of our iterative simulation model is that at the beginning of each model iteration we determine the current value of U_i^a for every agent a of type $i = 1, 2, 3$. If $U_i^a = 1$, then agent a does not move. If $U_i^a = 0$, then agent a moves according to the following rules:

1. Choose a random heading
2. Walk forward n steps, where n is a random number between 1 and 10.
3. Determine whether any agent occupies the new grid cell. If so, return to step 1.
4. If not, occupy the center of the new grid cell.

In Schelling's own models, the movement order was decided either by starting from the center and working outwards, or by starting from the upper left corner and sweeping right and down.⁶ In our models, we simulate a parallel process by randomizing the order of moves on each round. Also, Schelling had his agents move to the nearest empty square while our procedure does not always guarantee this. The model itself is iterated repeatedly until either an equilibrium state is reached (where every agent has achieved an MHP utility of one) or else a time limit of 1000 model iterations had elapsed.

For the analysis presented here, we consider instantiations of the model with community sizes ranging from $r = 1$ to $r = 24.5$ in increments of 0.5.

⁵Such utilities, first introduced by Simon (1971), distinguish only between *acceptable* and *unacceptable* alternatives.

⁶Schelling (1971, p.148)

This covers the range from just under the size of the Moore neighborhoods⁷ in Schelling’s model to just over the size of the entire virtual city. For each community size, we ran 100 repetitions, each starting from a different random configuration of the 498 agents on the grid. We also constructed a comparison set of 100 *random benchmark* configurations, each generated by a random initialization of our model without running the decision procedure.

For the purposes of later analysis, a coordinate file was recorded for each simulation that gives the location of each agent and its type. In all the analyses discussed in subsequent sections, we arbitrarily chose a single type of agent for analysis, and studied the extent to which this agent type exhibited either attraction to or segregation from the other two types.

3 Analysis of Segregation

Based on these simulations, our main finding is that, starting from random configurations, the equilibrium outcomes for most parameter values exhibit spatial segregation between agent types. The degree of segregation, however, depends crucially on the radius of vision r . When r is below 4.5, the size of segregated clusters are extremely small, and only exist as transients, not equilibrium states. Once we move to larger values of r , agents segregate themselves into increasingly larger clusters. It is only when r encompasses almost the whole population for each agent that segregation start to decrease, and eventually disappears when all agents are visible to each other. This suggests that in the present model, segregation is driven by agents having some, but not full, information about the wider community beyond their immediate neighborhood.

⁷To be more precise, a radius of $r = 1$ corresponds to a *rook neighborhood* (analogous to rook moves in chess) consisting of a cell together with the four cells sharing one of its faces. In these terms a Moore neighborhood is also called a *queen neighborhood*, since it includes the additional four cells sharing corners with the given cell.

3.1 A Scale-Sensitive Test of Segregation versus Attraction

Given the assumed symmetry among agent types, our analysis of segregation focuses only on a single type (Type 1), now designated as the *target population* of agents (labeled 1). Types 2 and 3 are then combined into a single *reference population* (labeled 0).⁸ In this context, the most widely used test of attraction versus repulsion (segregation) between such spatial point populations is based on cross K -function statistics⁹, which formed the natural starting point for our analysis. For each given community size, r , the cross K -function value, $K(d; r)$, denotes the expected number of reference individuals (per unit area) within distance d of a randomly selected target individual. These statistics are designed to detect either attraction or repulsion (segregation) between the target and reference populations. Roughly speaking, an unusually small (large) value of $K(d; r)$ is taken to indicate significant segregation (attraction) between these populations at *scale* d .

However, while tests based on cross K -function values did indicate the presence of significant segregation over a range of spatial scales, visual comparisons between typical equilibrium patterns and test results showed that these statistics sometimes failed to detect segregation that was readily apparent (especially for extreme values of r). Further analysis suggested that the main reason for this is that values, $K(d; r)$, focus *only* on the numbers of reference agents within distance d of a given target agent, and ignore the presence or absence of other *target* agents. Hence segregated clusters of similar agents were often missed.

This led us to consider a modified test statistic that incorporates counts of both types of agents. The particular form of this statistic (which is tailored to our present simulation model) is based on the observation that if the

⁸Note that since the MHP utility function U_a^1 for each agent, a , in the target population depends only on the *total size*, $(N_2^a + N_3^a)$, of the reference population in a 's current community, there is no need to distinguish between agent types within this reference population.

⁹ K -functions were first developed for univariate spatial point processes by Ripley (1976), and later generalized to *cross* K -functions for multivariate (marked) point processes by Hanisch and Stoyan (1979). Subsequent extensions include the influential paper by Lotwick and Silverman (1982), and are summarized in section 8.6 of Cressie (1993).

three identical populations of agents were randomly distributed throughout the virtual city, then in any given subregion, one would expect to find about the same numbers of each agent type. Hence, if we now let $N_0^a(d)$ denote the number of *reference agents* within distance d of given target agent, a , and let $N_1^a(d)$ denote the number other *target agents* within distance d of a , then under complete randomness one would expect $N_0^a(d)$ to be about twice as large as $N_1^a(d)$. Thus our test statistic is designed to estimate the expected difference, $E[N_0^a(d) - 2 N_1^a(d)]$, which should be about zero under complete randomness. However, to be precise, one must consider the *conditional expectation*, $E[N_0^a(d) - 2 N_1^a(d) | a]$, given the inclusion of target agent a . Here it is shown in the Appendix that if n_d denotes the number of distinct cells within distance d of any given cell, i.e. the d -neighborhood of this cell (which must be the same for all cells by our torus construction), then a correction term, $n_d/612$, is required in order to insure zero expectation with respect to every d -neighborhood (given our specific population sizes of each agent type). With this correction, the desired modification of cross K -functions for our present purposes is based on the *conditional expected local differences* defined for each agent, a , and distance, d , by¹⁰

$$\begin{aligned} \delta_a(d) &= E[N_0^a(d) - 2 N_1^a(d) - n_d/612 | a] \\ &= E[N_0^a(d) - 2 N_1^a(d) | a] - n_d/612 \end{aligned} \quad (2)$$

It is shown in the Appendix that under the null hypothesis of randomly located agents, $\delta_a(d) = 0$ for all d -neighborhoods.¹¹ Hence in each d -neighborhood of agent a , positive (resp., negative) values of $\delta_a(d)$ are associated with larger (resp., smaller) numbers of reference agents relative to

¹⁰As with cross K -functions, population counts in these local differences should in principle be normalized by population densities. But since such densities are constant, they have no effect on the tests developed, and are thus ignored for simplicity.

¹¹Technically, exact satisfaction of this zero-expectation property requires the additional simplifying assumption of binomially (rather than hypergeometrically) distributed population counts, as detailed in the Appendix. But for the present sample sizes, this approximation is very good. It should also be emphasized that this zero-expectation property is only for purposes of interpretation. As will be clear from the difference form of the final test statistic (8) to be used, the constant correction term, $n_d/612$, cancels out and has no influence on the tests to be conducted.

target agents than would be expected under randomness.

Observe next, from the identity of MHP utility functions for all agents and the complete symmetry of all cell locations on the torus, that this conditional expectation must be the *same* for all target agents. Hence the specification of a particular agent, a , in (2) could in principle be dropped. But for expositional purposes, it is most convenient to eliminate dependency on a by averaging. Thus we now focus on the *expected local difference* across all target agents, $a = 1, \dots, n_1 (= 166)$, at scale d , as defined by:

$$\delta(d) = \frac{1}{n_1} \sum_{a=1}^{n_1} \delta_a(d) \quad (3)$$

As a parallel to cross K -functions, these expected local differences can equivalently be interpreted as values of (2) for a randomly sampled target agent.

Before developing tests based on these indices, it should be noted that for d equal to the *community size*, r , there is a close relation between $\delta_a(r)$ in (2) and the values of a 's MHP utility in (1). In particular, since $N_0^a(d) - 2 N_1^a(d) > 0$ for each agent, a , *in equilibrium*, and since the correction term ($n_d/612$) tends to be relatively small by comparison,¹² one can expect to find positive values of $\delta(r)$ in equilibrium. So at equilibrium, one should not be surprised to find a substantial degree of attraction at the community scale, r , which is of course precisely what individual agents are trying to achieve. Thus our main interest focuses on the consequences of this behavior for segregation or attraction at scales, $d \neq r$.

With these definitions and informal observations, we can now formalize our testing procedure as follows. If at each community size, r , we let $\delta(d; r)$ denote the expected local difference in (3) for MHP processes under *community size*, r , [with corresponding conditional forms, $\delta_a(d; r)$, in (2)], and similarly, let $\delta(d; rand)$, denote the expected local difference in (3) for the *random benchmark process* defined above, then our testing procedure focuses on the *difference* between these mean values,

¹²Since $n_d \leq 35^2 - 1 = 1224$, it follows by definition that $0 \leq n_d/612 \leq 2$. This in turn implies that $E[N_0^a(d) - 2 N_1^a(d) \mid a] \leq \delta_a(d) \leq E[N_0^a(d) - 2 [N_1^a(d) + 1 \mid a]]$, and hence that the effect of this correction is small relative to the expected counts involved.

$$\Delta(d; r) = \delta(d; r) - \delta(d; rand) \quad (4)$$

at each scale, d , and community size, r . Given the above interpretation of expected local differences, positive (resp., negative) values of $\Delta(d; r)$ can also be taken to imply more attraction (resp., segregation) between agent types at scale d than would be expected under the random benchmark process. A test of the statistical significance of these Δ -values can thus be implemented in terms of a standard *difference-between-means* test.

To carry out such a test, we performed $N = 100$ simulations of the MHP process under each community size, r . If for each simulation, $s = 1, \dots, N$ we let $n_0^a(d; r, s)$ and $n_1^a(d; r, s)$ denote the observed values of $N_0^a(d)$ and $N_1^a(d)$ in that simulation, and let

$$\widehat{\delta}_a^{(s)}(d; r) = n_0^a(d; r, s) - 2 n_1^a(d; r, s) - n_d/612 \quad (5)$$

denote the corresponding (one-sample) estimate of $\delta_a(d; r)$, then the resulting estimate of $\delta(d; r)$ under community size r is given by

$$\widehat{\delta}^{(s)}(d; r) = \frac{1}{n_1} \sum_{a=1}^{n_1} \widehat{\delta}_a^{(s)}(d; r) \quad (6)$$

Similarly, for N parallel simulations of the random benchmark process, one can construct corresponding estimates

$$\widehat{\delta}^{(s)}(d; rand) = \frac{1}{n_1} \sum_{a=1}^{n_1} \widehat{\delta}_a^{(s)}(d; rand) \quad (7)$$

At this point it should be noted that the individual estimates, $\widehat{\delta}_a^{(s)}(d; r)$ [resp., $\widehat{\delta}_a^{(s)}(d; rand)$] are *not* independent within a given simulation (since d -neighborhoods can overlap one another). However, the estimates $\widehat{\delta}^{(s)}(d; r)$ [resp., $\widehat{\delta}^{(s)}(d; rand)$] for separate simulations, s , are independent by construction. Hence these can be regarded as N independent random samples of mean estimates from both the MHP process and random benchmark process.

Given this independence property, it follows that if the resulting sample estimate of $\Delta(d; r)$ for each simulation, s , is denoted by

$$\widehat{\Delta}^{(s)}(h; r) = \widehat{\delta}^{(s)}(d; r) - \widehat{\delta}^{(s)}(d; rand) \quad (8)$$

then the *grand mean* across all simulations, i.e.,

$$\widehat{\Delta}(d; r) = \frac{1}{N} \sum_{s=1}^N \widehat{\Delta}^{(s)}(h; r) \quad (9)$$

$$= \frac{1}{N} \sum_{s=1}^N \widehat{\delta}^{(s)}(d; r) - \frac{1}{N} \sum_{s=1}^N \widehat{\delta}^{(s)}(d; rand) \quad (10)$$

$$\equiv \widehat{\delta}(d; r) - \widehat{\delta}(d; rand) \quad (11)$$

yields an appropriate test statistic for discriminating between the MHP and random benchmark processes. In particular, this statistic satisfies the conditions for *Welch-Satterthwaite* (WS) difference-between-means test¹³ that allows for possibly different variances between these two statistical populations.

To carry out this test, we first calculated the sample means for the MHP process in (6) and the random benchmark process in (7) for each simulation. This was accomplished with a program (written in Matlab by one of the authors) that calculates $\widehat{\delta}^{(s)}(d; r)$ and $\widehat{\delta}^{(s)}(d; rand)$ at each scale, d , in each simulation, s . Here it should be noted that while community sizes, r , are allowed to range up to the maximum torus distance of approximately $r = 24$, we restrict the scale values, d , to lie in the range, $d = 1, \dots, 15$. The reason for this is that our torus approximation of d -neighborhoods on large planar grids breaks down for distances greater than half the width of the original (35×35) square, i.e., for $d \geq 17.4$.¹⁴ These computed sample means formed

¹³In addition to independence, this test requires that the two sample populations be normally distributed. While individual difference statistics, $\widehat{\delta}_a^{(s)}(d; r)$, are generally not normal, their sample means, $\widehat{\delta}^{(s)}(d; r)$, were confirmed (by Shapiro-Wilks tests) to be sufficiently normal for testing purposes. Normality in this case can also be justified on theoretical grounds from the well-known asymptotic normality properties of sample means for locally dependent data.

¹⁴In particular, d -neighborhoods larger than this are seen to overlap themselves on the

the basic inputs for the WS test conducted at each scale, $d = 1, \dots, 15$, and community size, $r = 1, \dots, 24$.

The test results were determined by calculating the t -values for $\widehat{\Delta}(d; r)$ and applying the WS approximate t -distributions for two-sided tests of the null hypotheses, $\Delta(d; r) = 0$ for each choice of r and d . For all practical purposes none of these t -distributions were distinguishable from the normal distribution at the given sample sizes (of $N = 100$ for each population). So our formal rule was basically to conclude *significant attraction* for t -values above 1.96 and *significant segregation* for t -values below -1.96 . But since the computed t -values ranged so far outside this interval, it proved to be more informative to plot these t -values directly, and simply indicate the critical interval $[-1.96, 1.96]$ on the graphs. This is the convention employed in Figures 2 and 3 below.

3.2 Test Results

To summarize our test results, we begin by observing that for all community sizes, $5 \leq r \leq 24$, the simulations always settled to an equilibrium state. The only exception is the extreme full-visibility case of $r = 24.5$, where by definition no equilibrium is possible (since *no* agent can be in a strict minority). Moreover, the test results showed that these equilibrium states for community sizes $5 \leq r \leq 24$ always involve significant segregation below a certain *critical scale*, d_r , depending on r . So even though agents are maximizing their MHP utilities (i.e., are in minorities less than one third of their current community populations), they nevertheless find themselves in segregated neighborhoods at all scales not exceeding the critical scale, d_r , for their radius of vision, r . This is illustrated in the right hand panel (b) of Figure 1 below for a community size of $r = 8$, with critical scale, $d_8 \approx 5.5$. Here the black dots denote locations of target agents (Type 1) and the locations of reference agents (Type 2 and Type 3) are denoted respectively by circles and triangles. (The large circles in the figure will be discussed below.) Here it is visually evident that each agent type exhibits clustering

torus, so that some cells start to appear more than once in the same neighborhood.

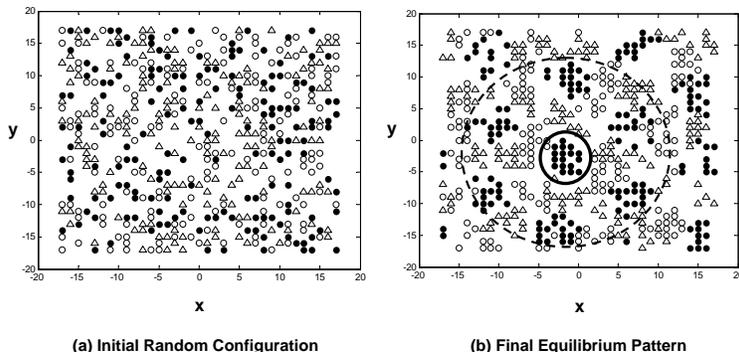


Figure 1: Example equilibrium pattern ($r=8$).

that results in segregation at certain scales, as will be verified by the test results below. For purposes of comparison, the left panel (a) of Figure 1 shows the initial realization of the random benchmark process leading to equilibrium pattern (b).

Turning now to the test results themselves, it is again useful to illustrate these findings for the case of community size, $r = 8$, in Figure 1(b) above. The full plot of t -values for all relevant scales, $d = 1, \dots, 15$, is shown for this case in Figure 2 below. Here the relatively narrow width of the critical interval, $[-1.96, 1.96]$ (indicated by the horizontal dashed lines in the figure)¹⁵ shows that there is a strong degree of statistical significance in both the regions of attraction and segregation in the figure. As for the shape of this plot, notice first that as predicted above (in the discussion of expected local differences), there is a peak of very *significant attraction* at precisely scale $d = 8 = r$, reflecting the MHP utility-maximizing behavior of all agents.¹⁶ More important for our present purposes is the lower end of the graph, where there is seen to be *significant segregation* between agent types at all scales, $d \leq d_s \approx 5.5$. (This will be illustrated more fully in Fig-

¹⁵As mentioned above the critical intervals for each of these 15 t -distributions are slightly different. But all are so close to the normal distribution, that such differences are imperceptible.

¹⁶As for the actual significance level in this case, the calculated t -value is seen from the figure to be more than 20 standard deviations from the mean. Thus the associated p -value is virtually zero, and indicates why t -value plots are more meaningful here.

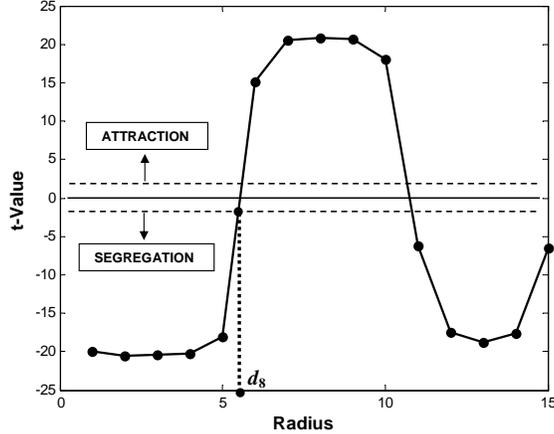


Figure 2: Plot of t -Values for Community Size 8

ure 4 below.) Here segregation appears to be most significant at a scale of about $d = 3$. Visually this corresponds to circular neighborhoods of radius 3 about each cell, as illustrated by the solid circle near the center of Figure 1(b). This particular neighborhood is seen to be just large enough to cover the cluster of target agents (black dots) shown in the figure. More generally, the most significant scales of segregation (below the critical scale, d_r) for each community size, $r \geq 5$, tend to reflect the cluster sizes of agent types seen in typical equilibrium patterns for that community size. (An alternative representation of such cluster-size relations is given in Figure 5 below.) Thus, while Figure 1(b) shows only one instance of a simulated equilibrium pattern for community size, $r = 8$, it helps to illustrate the broader statistical results summarized in Figure 2. Note finally that beyond the attraction peak at $d = 8 = r$, there appears to be a second peak of significant segregation at about $d = 13$. But further analysis shows that secondary peak is essentially a reflection of the regularly space clusters of agents at smaller scales. For example, the dashed circular neighborhood of radius 13 that is concentric with the smaller circle in Figure 1(b) is seen to be just large enough to include the second wave of target-agent clusters adjacent to the central cluster shown. Hence the concentration of target agents at this scale

is a consequence of similar concentration at smaller scales.

To analyze further properties of these statistical results, it is useful to distinguish three ranges of community sizes, r . The first is the *minimal range*, involving community sizes $1 \leq r < 5$, in which equilibria do not always occur. Second is the *intermediate range*, $5 \leq r \leq 20$, which covers the most interesting cases for our purposes, and third is the *maximal range*, $r > 20$, which exhibits behavior somewhat different from that of the intermediate range. Here we begin with the most important intermediate range, and then consider each of the extreme ranges in turn.

3.2.1 Intermediate Community Sizes

The key feature of the intermediate range, $5 \leq r \leq 20$, is the relation between each community size, r , and its associated critical scale, d_r . Specifically, as the radius of vision increases, the critical scale for segregation also increases. This is illustrated in Figure 3, where the t -values are plotted for selected values of intermediate radii, $(5, 8, 10, 12, 17, 20)$,¹⁷ and where each curve is numbered by its associated community size.

For sake of visual clarity, only the most relevant portion of each curve is shown, specifically the t -values at all scales, $d \leq r + 1$, for each community size, r . For example, the curve shown for community size, $r = 8$, is seen to be the lower portion of the same t -value curve in Figure 2. By definition, the critical scale, d_r , for each community size, r , is seen to be the scale at which the associated t -value curve first intersects the critical interval from below (shown by a black dot on each curve). So on curve $r = 8$, the black dot is seen to be at scale, $d_8 \approx 5.5$. Hence the steady increase in these critical scale values is easily seen by the progression of these intersection points.

Estimating the Number of Agent Clusters To obtain further statistical insight into this segregation phenomena, we constructed a program that attempts to identify the number of spatial clusters in the target population. This program involves two stages. First, the equilibrium configuration of

¹⁷These specific values were chosen in order to achieve a relatively even spacing between t -value curves.

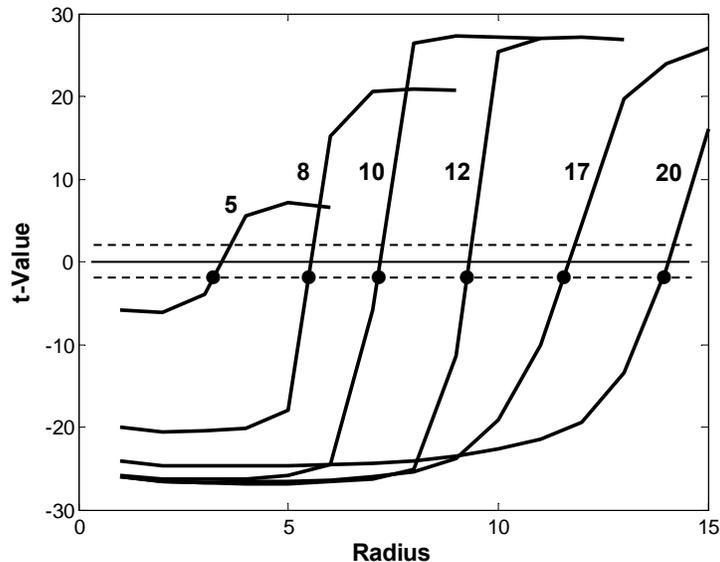


Figure 3: Community sizes 5 to 20

n individual target agents was successively grouped into smaller numbers of clusters by continually adding closest clusters (where distance between clusters was measured as the distance between their closest points). This hierarchical clustering procedure was continued until the n points corresponding to the target population were reduced to a specified cutoff number of clusters (which was set sufficiently low to include all clustering levels of potential interest).¹⁸

With a cluster hierarchy in hand, the second step was to determine the best number of clusters with which to group the target population. This number was not always obvious on visual inspection and exhibited substantial variation between simulations (with different initial random configurations of agents). After considerable experimentation, it was found that two criteria seemed to place reasonable bounds on the number clusters observed visually. The first criterion (which focuses on achieving compact group-

¹⁸This stage of the program was implemented using the Matlab package, clusterdata.m, with the ‘single linkage’ (or ‘nearest neighbor’) criterion for joining clusters.

ings within each cluster) is to find the cluster configuration that minimizes the mean squared distance of points to their respective cluster centroids.¹⁹ Since this criterion value, C_1 , for each cluster partition is always decreasing in the number of clusters, it was decided to compare observed values with those expected under random configurations. Here cluster hierarchies for 1000 random configurations were constructed and used to estimate the sampling distribution of C_1 under randomness. When observed criterion values, C_1 , are replaced with their z -scores, Z_1 , under this distribution, the cluster number with *minimum* Z_1 value can be interpreted as that for which C_1 is most significantly lower than would be expected under randomness. The second criterion (which focuses on achieving good separation between clusters) starts by computing for each cluster the shortest distance between points in the cluster and points outside, and then finds the cluster configuration that maximizes the mean of these shortest distances for each cluster. Since these values, C_2 , are also monotone decreasing in the number of clusters, the same random cluster hierarchies were used to construct a sampling distribution of C_2 under randomness. When observed values C_2 are replaced by their z -scores, Z_2 , the cluster number with *maximum* Z_2 value can be interpreted as that for which C_2 is most significantly higher than would be expected under randomness. As stated above, the optimal cluster numbers, N_1 and N_2 , produced by these two criteria tended to place reasonable bounds on the number of clusters observed visually over a wide range of examples. Hence our final representative *cluster number*, N was obtained by taking the average, $N = (N_1 + N_2)/2$, of these two optima.

To analyze these cluster numbers, we simulated 1000 equilibria for each possible community size, $r = 5, \dots, 24$, and constructed corresponding cluster numbers, N_{ir} for each simulation $i = 1, \dots, 1000$. But in order to interpret these results, it is important to note that cluster numbers by themselves are only meaningful relative to the particular size of the torus approximation used. Hence it was deemed to be more appropriate to convert each cluster number, N_{ir} , to a corresponding *average cluster size*, $S_{ir} = 166/N_{ir}$, in

¹⁹This is the implicit objective function employed in k-means clustering.

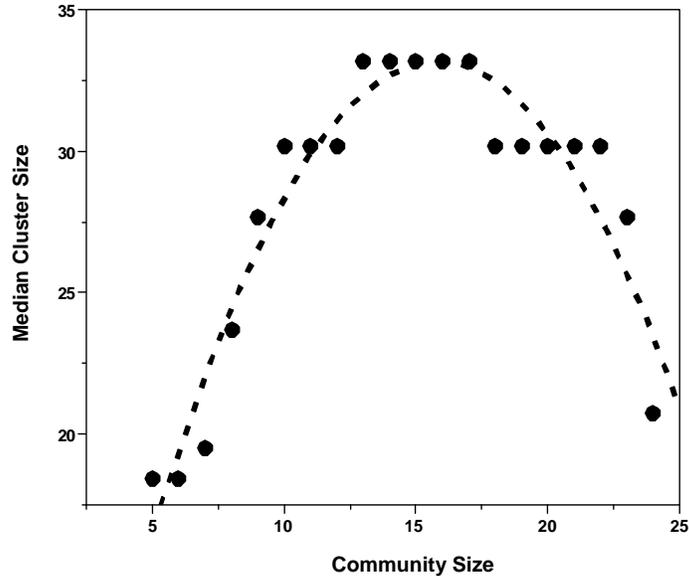


Figure 4: Equilibrium cluster sizes

terms of target agents per cluster.²⁰ In particular, the average size of these clusters relates more directly to the *degree of segregation* exhibited by the agents. In Figure 4 the median of these cluster sizes, $(S_{ir} : i = 1, \dots, 1000)$ is plotted for each community size, $r = 5, \dots, 24$. The step function nature of the plot is a consequence the underlying discreteness of possible cluster numbers. But even with this lack of smoothness, the regularity of its overall pattern is evident.²¹ In particular, median cluster size (degree of segregation) first increases with community size up to a radius of about $r = 16$, and then decreases with all larger community sizes (as emphasized by the simple quadratic fit shown as a dashed curve). This is in rough agreement

²⁰For example, if the present torus grid together with a given equilibrium cluster configuration were expanded to a square panel of nine identical copies, then by our torus construction, this would automatically constitute a larger equilibrium configuration for 9×498 agents. But while the number of target-agent clusters would increase nine fold, the *average cluster size* in terms of target agents would remain the same.

²¹While a plot of mean (rather than median) values tends to yield somewhat smoother results, the inherent step-function nature of this data is much clearer with medians.

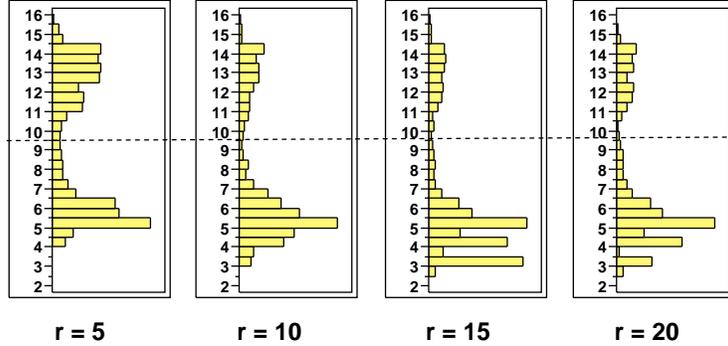


Figure 5: Cluster Size Distributions

with the segregation results above (in terms of critical scales), except for the slight decrease in cluster size from $r = 17$ to $r = 20$. But even here it should be borne in mind that this decrease from about 33 agents to 30 agents per cluster corresponds to the smallest possible increase in median cluster numbers from 5 to 5.5.²² So substantial decreases in cluster sizes are only seen to occur beyond $r = 20$.

It should also be noted that the cluster identification procedure underlying this plot is far less precise than the statistical significance results above. But in spite of its lack of precision, this procedure served to identify one additional feature of these equilibrium patterns that was entirely missed by our difference-between-means tests.

A Secondary Equilibrium Mode The histograms of cluster numbers identified at each community size revealed that there are actually *two* modes of equilibria: a *primary mode* and a *secondary mode*. In Figure 4 we have displayed these histograms for four community sizes, $r = 5, 10, 15, 20$, that are typical of the full range of sizes. Here the primary mode (small numbers of clusters) corresponds to the segregated equilibria that are the focus of this paper, and the secondary mode (large numbers of clusters) corresponds

²²Recall from the averaging definition of cluster numbers, $N = (N_1 + N_2)/2$, that the closest value above 5 is $(5 + 6)/2 = 5.5$.

to equilibria in which essentially no segregation is evident.²³ One should hasten to add that, except for the most extreme community sizes (close to either $r = 5$ or $r = 24$), the primary mode involves about 75% of all equilibria. But nonetheless, this secondary mode is of some interest.

First of all, while it is not difficult to construct rather uniform equilibrium configurations in which all agents are maximizing MHP utility, it was not clear to us that such equilibria could be even locally stable. However, these results show that many are. Moreover, the presence of such equilibria suggests that unless the initial random configuration involves some minimal degree of clustering (as generally seems to be the case), segregation may fail to emerge in equilibrium.²⁴ It is also of interest to note that unlike the primary mode, these non-segregated patterns appear to be very insensitive to community size. As shown by the dashed line in Figure 5, this mode always ranges between about 10 and 16 clusters with mean around 12.5. Hence, these dispersed patterns appear to exhibit roughly the same degree of agent heterogeneity at all scales.

Finally, in view of the non-segregated nature of this secondary mode, it is important to ask whether the inclusion of this mode might distort the relation of cluster sizes to community sizes occurring within the more important primary mode. This can be checked by including only cluster numbers, N_{ir} , within in the primary mode (i.e., $N_{ir} < 10$) and constructing a median plot paralleling Figure 4. Given the invariant shape of this secondary mode across community sizes, it is perhaps not surprising to find that the primary-mode plot is qualitatively almost exactly the same. Hence the only effect of this secondary mode is to inflate the level of median cluster numbers obtained. But since the ultimate effect of this inflation is to yield lower (more conservative) estimates of the degree of segregation present, we have chosen to show only the combined plot.

²³In fact, examples show that most of the clusters identified in this secondary mode are essentially an artifact of our clustering algorithm (which always identifies some number of clusters, even for very dispersed patterns).

²⁴The identification of such clustering thresholds is left for future research.

3.2.2 Minimal and Maximal Community Sizes

When agents' community size is very small ($r < 5$), the model fails to reach an equilibrium state in which all agents have maximal MHP utility. Nevertheless, the underlying dynamics of the model continue to exhibit non-random spatial structure that we set out to investigate.

We began with 498 agents of three types distributed randomly over our toroidal grid. Like the procedure used to study medium neighborhood sizes, we investigated 100 simulations for each neighborhood size corresponding to a different initial distribution of agents. Since the simulations never reach equilibria, we chose to sample the pattern reached after 1000 iterations of the model. This set of 100 samples was then compared to a set of 100 randomly generated distributions of agents on the grid using precisely the same procedure as above. In particular, we computed the appropriate $\widehat{\delta}^{(s)}$ values in (6) and (7) above for each simulation, $s = 1, \dots, 100$, and then tested for significant differences between them.

For these small neighborhoods, we found patterns that correspond to the secondary equilibria discussed above. Visually, this appears as very small clusters - more structured than completely random, but not statistically significant. These clusters are unstable, as the agents in them all have low utility and will continue to move seeking higher utility. Nevertheless, they persistently reappear as the model is iterated.

Finally, for large community sizes ($r > 20$) we found that equilibria continued to emerge all the way up to the maximal community size of $r = 24.5$ (in which no equilibrium is possible). However, the critical scales, d_r , for significant segregation in such equilibria began to diminish. This is also seen in Figure 4 above, where median cluster sizes diminish rapidly for $r > 20$. A possible explanation for this diminishing effect is given in the next section.

4 Discussion

Schelling's original result was striking, but perhaps not altogether surprising. While the degree of segregation observed in his model is far higher

than any individual's preference for homogeneity, agents still desired some degree of homogeneity in their neighborhoods. In our model, agents want to be surrounded by agents that are different from themselves. So why does segregation still arise in the model?

We believe that the result can be explained by careful consideration of the micro-dynamics of MHP models. Consider a model like ours with only two types of agents, and start with a population of only two agents (a and b) of the same type in close proximity. Because they have maximally heterogeneous preferences, these agents will be repelled from one another until they move outside one another's community. This repulsion, however, begins to be modified as additional agents are added to the model.

For example, suppose now that three new agents (c , d , and e) of the opposite type are placed in close proximity to the original two. Agents a and b will still repel one another, but will also both be attracted to c , d , and e because they are different in type. This means that a might now tolerate being close to b , because there is a chance to be in a strict minority relative to c , d , and e .

A more formal way of thinking about these dynamics is to consider the informational state of the agents as they make decisions. As each agent relies on exactly the same decision rule and utility function, only two factors can cause agents to make different movement decisions: their agent-type and their information about the locations and types of other agents, which we call the agent's *information set*.

Consider how these two factors interact. Each agent's information about their surroundings is strictly defined by their radius of vision. As such, this information set is position-dependent. A consequence of this position-dependence is that as two agents move closer together, their information sets become increasingly correlated. When they are adjacent, their information sets are very nearly identical. For two agents that are of different types, this near-identity of information sets will not necessarily lead to identical behavior, because their types will lead them to have different interests. However, for two agents that are of the same type, we should expect that their movement decisions to be highly correlated. Thus, correlations in information

sets will correspond to correlations in decisions about where agents want to move. As the relocation process continues, local groups of target agents will tend to move toward similar positions. This will in turn lead to further information convergence about the agents in their respective neighborhoods. In this way, similarity of information sets becomes its own attractor.

This attractive power of informational similarity is possible in our model because of the focus on *communities* rather than neighborhoods. Neighborhoods are fairly small, and can contain relatively few agents. However, depending on their radius of vision, the communities relevant for agents can be much larger geographically, and hence can contain many more agents. For example, in our present model suppose we let S_r denote the set of cells within distance r of any given cell, and consider vision radius, $r = 8$ (as discussed in Section 3.2 above). Then it is readily verified that the relevant community, S_8 , for any target agent, a , can contain as many as 197 agents. In this context, if we consider a neighborhood, S_4 , of radius 4 about agent a (which is less than the critical scale of $d_8 = 5.5$), then it can also be verified that S_4 can contain at most 49 agents. Hence even if all agents in neighborhood, S_4 , are also target agents, it is still quite possible that this population is a strict minority in a 's community, C_8 . Moreover, as vision radius r increases, the maximum-population size of C_r grows in proportion to r^2 .²⁵

More generally, considerations of both (i) neighborhoods versus communities and (ii) the overlap between agents' information sets offer possible explanations of the systematic changes in critical scales and median cluster sizes as vision radius changes. First for small radii, r , since the communities of agents are not much larger than their immediate neighborhoods, no equilibrium is possible in which target agents have even a few other target agents as neighbors. It is only when r becomes sufficiently large relative to local neighborhoods that sizable clusters of target agents are possible in

²⁵Here it is worth noting that this dominating influence of community over neighborhood will remain true even if the agents *discount* their interest in other agents based on distance separation — so long as discounting is not too severe. For example, if agents at distance r are discounted in a manner proportional to $r^{-\alpha}$, then the “effective” community population will continue to grow with r as long as $\alpha < 2$.

equilibrium. Moreover, as r increases, the overlap between the communities (information sets) of adjacent agents must also increase, so that adjacent agents are more likely to share minority positions in their similar communities. These two factors together thus help to account for the initial growth in both critical scale values and median cluster sizes as r increases from small values.

At the other extreme, recall that when r is maximal and all agents are visible to each other, no agent can be in a strict minority. In this extreme case target agents will be equally dissatisfied at all locations. In particular locations near target agents will be no more attractive than other locations. But as r decreases from this extreme, information sets will differ between locations, and in particular, will now be more correlated at locations close to one another. So as the system approaches equilibrium, with many target agents already satisfied, it is more likely that locations close them will also yield satisfaction for relocating target agents. So again one can expect to see larger clusters of target agents appearing in equilibrium as r decreases from its maximal value.

5 Conclusions

We set out to investigate the robustness of Schelling's famous segregation result, subjecting it to the most stringent test that we know. Our model shows that MHP utility functions, where agents want to be in the smallest minority, can themselves create segregated clusters. So long as agents have only partial information about their entire city, and consider their community to be a few blocks in size, we find robust segregation in equilibrium.

This result, combined with those previously established, show the overwhelming prevalence of segregation when agents base their decisions on racial preferences. So long as agents all base their decisions on the same property, regardless of their attitudes towards it, we will tend to find segregation in a population. In this model, we aimed to provide the strongest possible test of this hypothesis, by having agents actively seek out those that are different than themselves. However, by adding the very weak, re-

alistic assumption that people care about their community rather than just their immediate neighbors, we find that individuals still typically end up in a segregated state.

This result tells us two things: First, it tells us that Schelling's original model is extremely robust to various kinds of perturbation. Second, that because of this robustness, we have a single macrobehavior that could be supported by a wide variety of underlying micromotives. Let us consider each in turn.

The robustness of Schelling's original model can be best explained by the fact that the underlying dynamic relies only on a tipping phenomenon. So long as all the agents have the same structure to their preferences, initial innocuous movements will quickly cascade across the entire population, and we have segregation as an absorbing state of the model. In our model, we saw that informational similarities were sufficient for this cascade to occur. Informational similarities result in highly correlated movement decisions, and as agents move more toward each other, their information sets become increasingly similar. This is a rather different dynamic from the standard Schelling model. In the original model, we get cascades of movement because agents were weakly repelled by agents of a different type, and once enough of their fellows have moved out of their neighborhood, their threshold for happiness is no longer met, and they move themselves. This creates a separating equilibrium, in which each type is most easily satisfied by being amongst its own. Nothing in that model relies on informational similarities. But in each model, a cascade of segregation occurs and quickly takes over the population.

This robustness reveals why Schelling's initial insight is so essential. We cannot argue from the observation of a macrobehavior to an understanding its micro-foundations without a great deal of additional investigation. While determining whether a model's results are robust or not is an essential component to most model validation, robust models can pose their own unique scientific challenges. When we have a model with very fragile results, where the phenomenon under study only manifests itself under quite limited conditions, we can often be in a superior epistemic position. If the fragile

model's assumptions are appropriately realistic and calibrated against empirical information, then we have a much better chance of having captured the precise underlying dynamics for the macrophenomenon in question. This suggests that intervention strategies can be more easily studied and implemented. However, with robust modeling results the same phenomenon can be overdetermined by competing explanations. As such, we learn less about the empirical phenomenon than we do in the fragile cases; we cannot identify what micro-dynamic is correct without further study.

Although segregation is extremely robust in the family of Schelling-like models, not all cities are massively segregated and perhaps some cities are well-integrated. This leads to the unsurprising observation that these models do not capture the precise population dynamics unfolding in cities, which is unsurprising given their simplicity. However, they remain important tools both for fundamental understanding of population dynamics, but also for thinking about the effects of potential policy interventions.

Segregation's robustness suggests that most policy interventions are likely to fail to prevent one of the many possible micro-dynamics from taking over and continuing to drive segregation. But this prediction is testable. If we can find and study full or partially integrated communities, then we would discover that the Schelling model and its variations fail to capture some feature of the real world that is essential for a proper understanding of the segregation phenomenon. So while we have seen that the original Schelling model is quite robust, it is possible that it is too robust. Given its strong predictions of the near-certainty of segregation, this robustness makes it easy to mount challenges to the completeness of the model.

References

- [1] Bruch, E. and Mare, R. (2006). Neighborhood Choice and Neighborhood Change, *American Journal of Sociology*, 112(3):667-709.
- [2] Cressie, N. (1993). *Statistics for Spatial Data, revised edition*. New York: Wiley

- [3] Fossett, M. and Dietrich, D. (2009). Effects of city size, shape, and form, and neighborhood size and shape in agent-based models of residential segregation: are Schelling-style preference effects robust? *Environment and Planning B: Planning and Design*, 36(1):149-169.
- [4] Hanisch, K-H and Stoyan, D. (1979). Formulas for the second-order analysis of marked point processes. *Mathematische Operationsforschung und Statistik Series Statistics*, 10, 555-560.
- [5] Lotwick, H. W. and Silverman, B. W. (1982). Methods for analysing spatial processes of several types of points. *Journal of the Royal Statistical Society, Series B*, 44, 406-413.
- [6] Pans, R. and Vriend, N. (2007). Schelling's spatial proximity model of segregation revisited. *Journal of Public Economics*, 91(1-2):1-24.
- [7] Ripley, B. D. (1976). The second-order analysis of stationary point processes. *Journal of Applied Probability*, 13, 255-266.
- [8] Simon, H. A. (1957). *Models of Man: Social and Rational*. New York: Wiley
- [9] Schelling, T. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, 1:143-186.
- [10] Zhang, J. (2004). Residential segregation in an all-integrationist world. *Journal of Economic Behavior & Organization*, 54(4):533-550

Appendix

In this Appendix we derive the normalization term required for the zero-expectation condition that $\delta_a(d) = 0$ hold for all d -neighborhoods. In particular, we calculate the conditional expectations of both $N_0^a(d)$ and $N_1^a(d)$ given the presence of agent, a , and use these to normalize the expected difference, $N_0^a(d) - 2 N_1^a(d)$. To do so, we begin by observing that the null hypothesis of random assignments of each agent type to distinct cells on the torus grid of 1225 ($= 35^2$) cells yields a standard “urn sampling” problem with associated hypergeometric distributions for the random variables, $N_0^a(d)$ and $N_1^a(d)$. But given the population sizes of 166 agents of each type, these distributions are well approximated by binomial distributions (which technically involves sampling *with replacement*, where more than one agent

can in principle occupy the same cell). Given this approximation, and recalling from the text that n_d denotes the number of distinct cells within distance d of the cell occupied by agent a (which is the same for all cells), the conditional expectation of $N_0^a(d)$ given the presence of a is simply

$$E[N_0^a(d)|a] = n_d p_0$$

where $p_0 = (2 \cdot 166)/(35^2 - 1) = 332/1224$ is the probability that any given cell other than a 's cell is occupied by a reference (type 0) agent. Similarly the conditional expectation of $N_1^a(d)$ given the presence of a is of the form

$$E[N_1^a(d)|a] = n_d p_1$$

where $p_1 = (165)/(35^2 - 1) = 165/1224$ is the probability that any given cell other than a 's cell is occupied by one of the other 165 target (type 1) agents. Thus it follows at once that

$$\begin{aligned} E[N_0^a(d) - 2 N_1^a(d)|a] &= n_d p_0 - 2n_d p_1 \\ &= n_d (p_0 - 2p_1) \\ &= n_d (332 - 2 \cdot 165)/1224 \\ &= n_d/612 \end{aligned}$$

and hence from (3) that for all d ,

$$\delta_a(d) = E[N_0^a(d) - 2 N_1^a(d)|a] - n_d/612 = 0$$

must hold under these assumptions.