Article

Environment and Planning A

Environment and Planning A
2016, Vol. 48(10) 2046–2066

DOI: 10.1177/0308518X16653402
epn.sagepub.com

How diverse can measures of segregation be? Results from Monte Carlo simulations of an agent-based model*

Daniel Arribas-Bel
Department of Geography and Planning, University of Liverpool, Liverpool, UK

Peter Nijkamp
Department of Spatial Economics, VU University Amsterdam, Amsterdam, The Netherlands

Jacques Poot
National Institute of Demographic and Economic Analysis, University of Waikato, Hamilton, New Zealand

Abstract

Cultural diversity is a complex and multi-faceted concept. Commonly used quantitative measures of the spatial distribution of culturally defined groups—such as segregation, isolation or concentration indexes—have been designed to capture just one feature of this distribution. The strengths and weaknesses of such measures under varying demographic, geographic and behavioral conditions can only be comprehensively assessed empirically. This has been rarely done in the case of multigroup cultural diversity. This paper aims to fill this gap and to provide evidence on the empirical properties of various segregation indexes by means of Monte Carlo replications of agent-based modelling simulations under widely varying assumptions. Schelling's classical segregation model is used as the theoretical engine to generate patterns of spatial clustering. The data inputs include the initial population, the assumed geography, the number and shares of various cultural groups, and their preferences with respect to co-location. Our Monte Carlo replications of agent-based modelling data generating process produces output maps that enable us to assess the sensitivity of the various measures of segregation to parameter assumptions by means of response surface analysis. We find that, as our simulated city becomes more diverse, stable residential location equilibria require the preference for co-location with one's own group to be not much more than the group share of the smallest demographic minority. When equilibria exist, the values of the various segregation measures are strongly dependent on the composition of the population across cultural groups, the assumed preferences and the assumed geography. Index values are generally non-decreasing in increasing preference for within-group co-location. More diverse populations yield—for given preferences and geography—a greater degree of spatial clustering. The sensitivity of segregation

*Additional results as well as working code and earlier working paper versions are available at http://darribas.org/

Corresponding author:
Daniel Arribas-Bel, Department of Geography and Planning, University of Liverpool, Liverpool, UK.
Email: D.Arribas-Bel@liverpool.ac.uk

measures to underlying conditions suggests that meaningful analysis of the impact of segregation requires spatial panel data modelling.

Keywords
Cultural diversity, segregation, agent-based model, Monte Carlo simulation

Introduction

International migration is not just a demographic phenomenon of population movement from one country to another. It also leads to both quantitative and qualitative social and economic compositional changes and related impacts in new host regions, while it also impacts on the interaction patterns between people either locally (e.g., between natives and immigrants) or internationally (e.g., between sending and receiving countries). A noteworthy phenomenon is that migrants exert, once they have settled in a given destination area, centripetal or centrifugal effects on others in the same area—either with respect to people with the same cultural identity, other migrant groups, or local natives. Consequently, the social geography of a city may exhibit drastic changes in its socio-cultural composition and spatial distribution of culturally defined groups over several decades after an initial migration influx. Cultural diversity triggers a range of spatial socio-economic phenomena related to housing and labor markets, social capital, social networking, bridging and bonding (see also Brickell and Datta, 2011; Nijkamp et al., 2012; Samers, 2009; Simon, 1999). Cultural diversity, as a result of a migration influx, may function either as an attraction or a repulsion force. In many cases, it appears to act as a driver for spatial movements of people who like or dislike a certain socio-cultural composition in a given urban area (Urry, 2000) and hence are inclined to move to or from this area. This externality phenomenon has been extensively discussed and analyzed by means of the so-called Schelling model (Schelling, 1969, 1971, 1978).

The core driver of Schelling's compositional dynamics model is the anticipated utility derived from population composition in urban neighborhoods. Cultural diversity is—beyond a critical threshold—a key factor in location decisions of people. In its broadest form, diversity may refer to religion, ethnicity, language, gender, age, income, life style, profession, etc. and the potential interactions between such population characteristics. Observed cultural diversity is closely connected to social cohesion attitudes or behaviors of residents in a city, whose location is often interpreted as the outcome of a process of minimizing or resisting so-called residential stress (Brown and Moore, 1970). If this cultural diversity is related to a wealth of differences across a wide range of nationalities (especially in modern mega-cities such as London or New York), one may even speak of super diversity (Vertovec, 2007).

The resulting spatial movement patterns as a result of either cohesion or dissimilation may create complex urban dynamics. After a move of one person or household with specific features, the cultural diversity of an area will only marginally change, but this move may prompt a chain reaction of other people in the same neighborhood. The simulation of such dynamic processes can be either based on microsimulation models or agent-based models.[1] Especially, the latter class of models has become fashionable over the past decade and has prompted a great variety of micro-based studies on dynamic human behavior, often including the question whether a new social and spatial equilibrium may emerge after a small initial disturbance. Such models are also increasingly used in resilience studies.

In this paper, we are concerned with measuring the degree of diversity in the cultural make-up of sub-areas of a city or region. However, common measures—such as segregation, isolation or concentration indexes—have been designed to capture just one aspect of the spatial distribution of the population across culturally defined groups. The strengths or weaknesses of such measures under varying demographic, geographic and behavioral conditions can only be comprehensively assessed empirically. To date, this has been rarely done in the literature in the case of multigroup cultural diversity. The present study aims to fill this gap. We use the classical agent-based model (ABM) from Schelling (1971) as the "theoretical engine" to generate widely varying patterns of segregation. Inputs into each simulation include the initial population, the assumed geography, the number and shares of various cultural groups, and their preferences with respect to co-location. We provide evidence on the empirical properties of the various measures by using Monte Carlo replications of each ABM (MC-ABM) simulation. Our data generating process produces output maps that enable us to assess the sensitivity of the various segregation measures to parameter assumptions by means of response surface analysis. The combination of agent-based modelling with MC simulations, two usually separately employed techniques, together with the application of this approach to the evaluation of segregation measures in the presence of varying multigroup compositions of the population, are the main contributions of this paper.

We find, not unexpectedly, that for a given geographical configuration of neighborhoods certain population distributions and cultural preferences with respect to co-location are more compatible with equilibrium in residential choices than others. We also find that, as our simulated city becomes more diverse, stable residential location equilibria require the preference for co-location with one's own group to be bounded by the group share of the smallest demographic minority. When equilibria exist, the values of the various segregation measures are strongly dependent on the composition of the population across cultural groups. More diverse populations yield—for given preferences and geography—a greater degree of spatial clustering. Our measures of segregation are shown for our assumed stylized rectangular city geography to be non-decreasing in increasing preference for within-group co-location. However, for more realistic geographies such monotonicity is not assured.[2] The sensitivity of segregation measures to underlying conditions suggests that meaningful analysis of the impact of segregation requires spatial panel data modelling. Specifically, given that cities vary widely in geography and demographic composition, a certain value of a segregation index may be observed in cities with widely varying circumstances. The assessment of an impact of a change in segregation on measures of social and economic outcomes can then not be meaningfully done in a cross-sectional setting. Instead, panel data may permit a fruitful analysis with city fixed effects and potentially also spatial spillovers.

This paper is organized as follows. The next section is devoted to a concise introduction of our selected indicators of segregation that form the foundation for our empirical work, further details given in Appendix 1. Section 'The Schelling model as the theoretical engine' highlights the principles of the methodology employed in this study, with a particular emphasis on the classical Schelling model (reviewed in Appendix 2). Next, Section 'A Monte Carlo based approach' offers a concise description of the simulation design and strategy which produces the main results. These results are presented in section 'Results'. Section 'Concluding Remarks' concludes.

Measuring cultural diversity across space

Culture is a complex socio-anthropological phenomenon that concerns certain common behavior, perceptions and beliefs within groups of people that are defined by characteristics such as country of birth, ethnicity, race, ancestry or language (e.g., Kroeber and Kluckhohn, 1952). Cultural diversity can be either defined as the extent to which different cultural groups are present in society or, alternatively, the extent to which the distribution of different groups across space is dissimilar.[3] Cultural diversity of society does not necessarily imply cultural diversity across space and vice versa. A society may be made up of many ethnic groups but when the relative shares of each group are the same at every location there is no spatial dissimilarity or clustering. On the other hand, a society with just two distinct cultural groups that have different shares of the population at various locations is spatially diverse but not necessarily a culturally diverse society.

In this paper, we are concerned with the spatial kind of cultural diversity, i.e. the extent to which the cultural make-up of the population varies across pre-defined areas. When there is spatial variation in the cultural composition of the population, this is commonly referred to as spatial dissimilarity or segregation. It is usually measured in terms of residential location but can also be measured in terms of workplace location or other location-referenced activities. Traditionally much of the literature has been concerned with segregation in the presence of only two groups, e.g. "black" versus "white" in the USA, and the commonly utilized segregation measures reflected this (see Massey and Denton, 1988). More recently, the reality of populations being made up of many distinct cultural groups has led to multigroup segregation measures (see Reardon and Firebaugh, 2002). Additionally, researchers have become increasingly aware of the Modifiable Arial Unit Problem (MAUP) which results from the fact that the spatial units across which segregation measures are calculated are often administratively determined (e.g. census tracts) rather than based on actual human behavior and interactions. Varying the boundaries for given geography and behavior may lead to rather different segregation index values. For this reason, spatial segregation measures have been proposed that take the geo-referenced location of individuals, and distances between them, explicitly into account (see Reardon and O'Sullivan, 2004).

In general, segregation measures can be *local-non-spatial* (i.e., the extent to which the population share of a particular group at a particular location is different from its share in the total population) or *global-non-spatial* (the extent to which a group is distributed differently across areas as compared with the population overall). They can also explicitly take *location* and *proximity* into account. In this case, there can be *local-spatial* measures of dissimilarity (e.g., the extent to which the share of a group in a particular area differs from that of neighboring areas) or *global-spatial* measures (e.g., the correlation between a group's share in an area and the group's share in surrounding areas). In the present paper, we are concerned with global summary measures that describe the dissimilarity of the distribution of two or more groups across predefined areas, without taking distance-defined spatial correlation into account.

We select four widely used global non-spatial measures of segregation: the segregation index (SI_g), the Isolation index (II_g), the Ellison & Glaeser index (EG_g), and Theil's information index (TD). These are representative of many other similar indexes and fairly distinctive among themselves, ensuring a good coverage of different properties and perspectives. Appendix 1 presents each of them in full detail and discusses their intuition as well as some of their core characteristics. For an extensive review of group and area diversity measures, the reader is directed to Nijkamp and Poot (2015).

The Schelling model as the theoretical engine

In order to examine the behavior of the various indicators of segregation under a range of assumptions, our approach relies on computer simulations. We evaluate the indicators using observations generated under synthetic conditions. This allows us to control the data generating process and establish a workbench on which comparisons are appropriate and meaningful. The tool chosen to generate the maps of the spatial distribution of a culturally diverse population is the original segregation model of Schelling (1969, 1971, 1978), credited as one of the first ABMs, and the first one to tackle patterns of segregation. This is a robust and well-understood model that has been used by researchers for decades and thus represents a good theoretical engine to form the basis of our simulations. An additional advantage of using the Schelling model is that, give its simplicity and level of abstraction, it can be used to study segregation processes in a variety of contexts. Although we will use the notions of neighborhood and residents since this is one of the most common cases in the literature of studying segregation, the *containers* in which agents interact can be potentially interpreted as regions, workplaces or other entities to which agents may belong. A brief review of the Schelling model's main intuition as well as an overview of its use in the literature, particularly focusing on the latest related contributions, can be found in Appendix 2.

A few parameters in the model are key in determining the outcome. These include the neighborhood structure, the size and composition of the population, and the individual residential preferences. In our experiments, we modify all of these to cover a wide range of potential situations. The choice of neighborhood structure is among the first to be made when designing a Schelling-type model. Boundaries must be defined that create the neighborhoods within which agents will locate and interact. For the sake of simplicity, and to make results more comparable to most of the literature in this regard, we consider two alternative setups with different population distributions.[4] The one we term "dense" is composed of 25 neighborhoods with 200 locations in each of them; a sparser one we name "sprawl" contains 100 neighborhoods with only 50 locations in each of them. Both setups have thus a total of 5000 potential locations. For each of them, we consider two population sizes: 3750 and 4250. These result from applying a 25% and 15% vacancy rate, respectively. This means that, at any given point, there is a proportion of all the potential locations in the grid that is vacant. This ensures that agents can move around more easily and find locations where they are content to stay. Hence, equilibrium is achieved faster, as long as it is feasible.

The composition of the population is defined by two parameters: the number of groups and the shares they represent. These two are key instruments in determining the outcome patterns in our multi-group setup. When deciding all the cases to consider, it is important to balance the trade-off between obtaining a wide range of situations and to keep the number of cases small enough so that the computational burden remains bearable. With this compromise in mind, we carefully select six combinations of numbers and shares of groups to which the population is allocated. They are the following: a benchmark of two groups with equal shares (50%–50%); two groups with demographic majority and minority populations (70%–30%), which we will call the "single minority"; four groups with two symmetric majorities (40%–40%) and two symmetric minorities (10%–10%), which we will refer to as "duopoly with minorities"; four groups, with one clear majority (70%) and three symmetric minorities (10%–10%–10%), codenamed the "multiple minority"; four groups with a consecutively diminishing presence (40%–30%–20%–10%), referred to as the "ladder"; and the "diversity" case with five groups of equal weight (20% each).

In the simple model, we are using the preferences of an agent are reduced to one dimension: the *minimum* fraction of all individuals in the same neighborhood that must be of the same group before he or she is content. This fraction between 0 and 1 is represented by the parameter τ. The lower bound implies the agent is always content and the upper bound implies that he or she is satisfied only when everyone else in the neighborhood belongs to the same group. In any iteration dissatisfied agents are randomly reallocated to another neighborhood. In equilibrium, all agents in all neighborhoods are content and no longer wish to move. One of the main contributions of Schelling was to show that even very low values of τ can lead to important degrees of residential segregation. A higher preference for co-location implies that it is harder for the model to converge into a solution in which every agent is content.[5] We begin every simulation with $\tau = 0$, so that even a completely random allocation of agents represents a compatible outcome (and the initial situation is therefore immediately an equilibrium). This provides a benchmark from which we gradually increase the value of τ, forcing the system into more restrictive outcomes. In these cases, agents require larger and larger proportions of neighbors from their own group in order not to leave the neighborhood. We continue this process of increasing τ until *none* of the 500 draws (see section 4) we run for each combination of cultural groups, group shares and τ has converged. At that moment, that particular case is considered finished.[6]

We call each combination of vacancy rate, neighborhood and population structures, and τ a *scenario*. Combining every option along each dimension yields a total of 720 scenarios, although we only end up considering 326 due to the stopping logic explained above. Once the model has been run for a given scenario, it is possible to visualize the outcomes on a map. Figure 1 displays two outcomes of the model in two scenarios that differ only in the value assigned to the preference parameter. Both maps represent a dense geography with a vacancy rate of 25% and two groups, dark and light grey, who are each assigned half of the population. On the left side, individuals do not have a preference for the characteristics of their neighbors and, consequently, any outcome represents an equilibrium distribution. The pattern in this case is that of a random spread with colors not showing any clear spatial clustering. The right panel displays a similar setup but, in this case, individuals have a strong preference for similar neighbors. Specifically, *both* groups want at least 50% of the neighborhood population to be the same as them. This translates in a clearly distinct spatial pattern with complete segregation: agents are surrounded by other individuals in the same group.

Although illustrative, outcome maps such as Figure 1 are not the main aim of this paper. Our ultimate goal is to study summary measures of the inherent patterns in the maps. This means these maps become the *input*, rather than the output of the analysis. The ABM-generated data and maps are used to calculate the four diversity indexes listed in section 2 and formally defined in Appendix 1. Continuing the illustration, Table 1 shows the measures

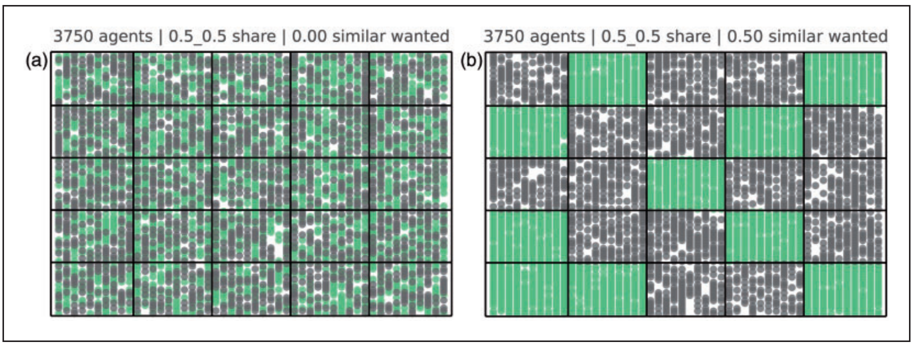

Figure 1. Output map examples (two groups): (a) $\tau = 0$ and (b) $\tau = 0.5$.

Table 1. Segregation indexes.

	Figure 1(a)		Figure 1(b)	
	Group 1	Group 2	Group 1	Group 2
Ellison & Glaeser Concentration	−0.000157	−0.000157	0.041355	0.041355
Isolation	1.017955	1.017955	2	2
Segregation	0.077333	0.077333	1	1
Theil	0.006505		1	

calculated for the two maps displayed in Figure 1. These have been computed for each group, using the counts of agents across each of the 25 neighborhoods and 5000 sites. It can be clearly seen that all indexes mirror the stark differences across maps: values in the left panel are low, reflecting the randomness of the distribution; while they are significantly higher on the right panel, as one would expect from a much more concentrated pattern.[7] This is the *output* that we analyze and base our conclusions on. This requires us to create a large sample of such results to draw sensible conclusions. For that reason, we turn to a Monte-Carlo approach, the other key computational tool in this study.

A Monte-Carlo agent-based approach

In order to compare the properties of different indexes of spatial segregation in a synthetic but statistical context, a handful of solutions to the Schelling model are not sufficient. We need a larger, more robust, set of solutions that can shed light on general patterns. Consequently, our next step is to scale up the approach from the previous section with two main objectives in mind: one, to eliminate the effect that randomness or unlikely initial conditions may have on the conclusions; and, two, to cover an extensive range of cases and situations that may lead us to general conclusions. The tool that allows us to meet these objectives is a MC approach. This widely used simulation method, combined with the ABM used at the core of this strategy, is instrumental to overcoming the main two challenges described above. In MC simulation, each scenario defined by the neighborhood structure, size and structure of the population, and strength of preferences for similar neighbors. Every scenario is run 500 times. This ensures that extreme initial values or particular solutions do not influence the final results. Each MC draw produces an outcome map like those in Figure 1, which is then used to calculate the segregation indexes listed in section 2 and formally defined in Appendix 1. Average values and the variance across replications are then represented in an organized way that permits direct visual comparison, not only within but also across different scenarios.

Everyone of the simulated scenarios is defined by a combination of five parameters: type of geography (“dense” or “sprawl”), vacancy rate (25% or 15%, in turn determining the total population size), preferences of agents (τ), number of groups, and share each of these represent. We run the model 500 times for each scenario and obtain a set of segregation indexes for each run that converges. This produces a substantial amount of output, which is very hard to explore and understand in raw form.[8] Our strategy to process these results is twofold. First, we generate a figure that contains one plot per index. Each of these graphs displays different values of τ (preferences) along the x axis, and the *average* value plus or minus three *standard deviations* of the index values across converging replications on the *y* axis. Second, we perform a response surface analysis (Box and Wilson, 1951) of the entire set of results by running regressions in which we model the value of each index obtained in each replication as a function of the parameters of the scenario for which it was calculated. This allows us to draw more systematic conclusions about the performance of the indicators and the effect of the scenario components on them.

Results

We generate up to 500 independent sets of results for every value of τ we test. The starting value of τ is always 0, increases by 0.024 and stops when the model no longer converges in any of the draws for a given τ, or when $\tau = 0.?7$ (the latter case yields only equilibria in the "benchmark" scenarios). As mentioned above, the chance of a model converging is related to the feasibility of equilibrium for a particular case. On one side of the spectrum, when $\tau = 0$, every initial situation is a solution, so the probability of the model converging is one. On the other extreme, when τ is high, implying a large proportion of neighbors are required to be alike for an individual to be content, there are scenarios in which equilibrium is impossible and the probability of convergence is thus zero. At the same time, the more demanding a set of parameters, the more iterations are needed to find a solution, if a solution does exist. Figure 2 displays, for the case of a "dense" geography and 25% of vacancy rate, the number of converging runs out of 500 and the average number of iterations—among the converging ones—that the model needed to reach equilibrium (topping at 2000, which is the limit for convergence after which, for computational reasons, we skip the scenario).[9] As can be quickly seen, the more groups there are and the smaller a minority is demographically, the more difficult it is to find a solution and, hence, the less restrictive preferences must be to reach equilibrium. This is a reflection in the model of a fairly realistic phenomenon: as societies become more diverse, stable residential location equilibria are only possible when people become more tolerant in general and the minorities in particular. More specifically, Figure 2 shows that stable residential location equilibria require the preference for co-location with one's own group to be not much more than the group share of the smallest minority. Although intuitively obvious, it is comforting to find a clear counterpart in our simulations arising from such a simple set of assumptions.

As mentioned above, the amount of information generated by simulations represents a challenge to organize and visualize. Our first approach to explore the behavior of the indexes is to structure the results in a set of charts arranged in a familiar way that allows for meaningful comparison. Figure 3 contains the results for the Theil index *TD* in a "dense" geography with a 25% vacancy rate, while similar plots for every other scenario and index can be found in the online companion to this paper. The figure is composed of six subplots, one for every population structure considered. These graphs display the average value and variance (the shaded area covers plus/minus three standard deviations) of the index over the at most 500 MC draws.[10]

We first note that Figure 3 shows, like Figure 2, that different population structures—and specifically the share of the smallest minority—determine the extent of tolerance required to make equilibrium feasible. Figure 3 also shows that the mean value of the index is increasing in τ (this is also the case for the other indexes, as can be seen from the figures in the online appendix). As one would expect, as preferences become more and more inclined toward similar neighbors, the outcomes become more spatially clustered, here signaled by the Theil index *TD*. However, the functional relationship between τ and *TD* depends strongly on the structure of the population. The values start always at zero when preferences are

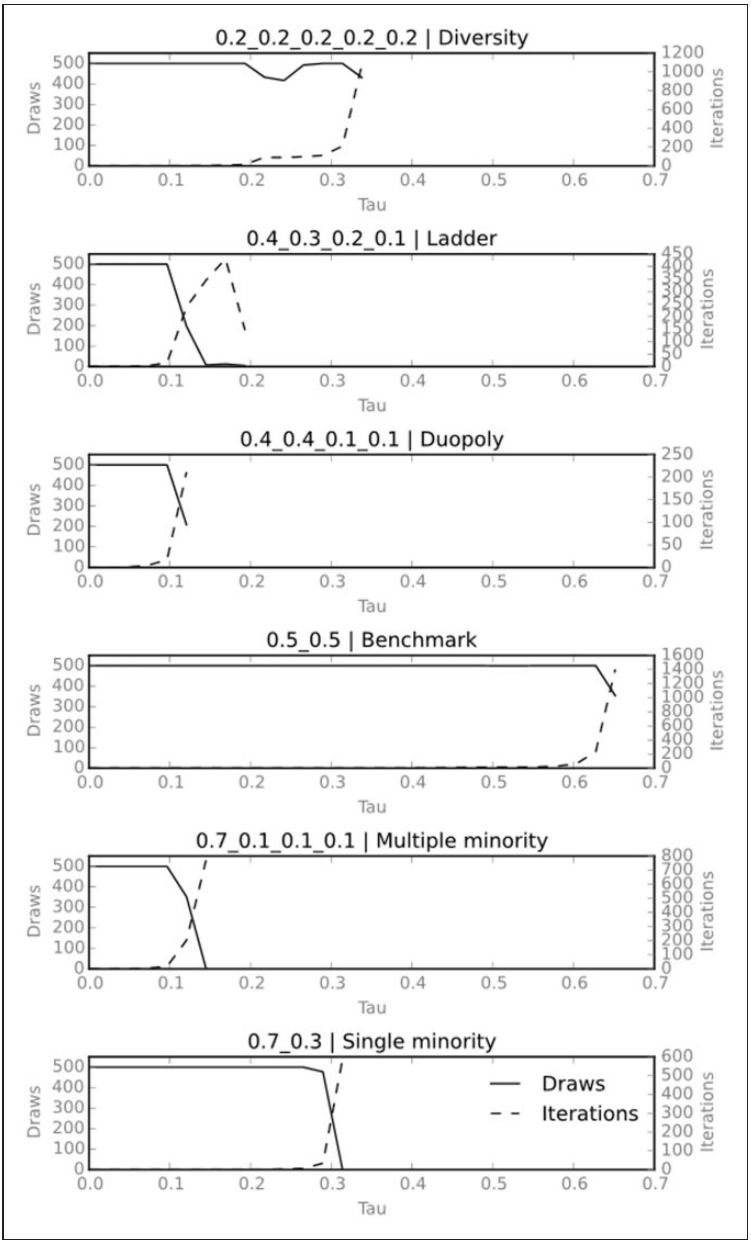

Figure 2. Simulation descriptives for the case of a "dense" geography and 25% vacancy rate.

non-constraining, but take off not far below a τ value equal to the share of the smallest minority. Equally, the progression of the index as τ increases differs across plots. In some cases, such as the duopoly with minorities, it is close to a linear relationship, while in others, such as the benchmark, it is best thought of as an *S*-shaped function. Given the sheer amount of combinations and the space constraints, we refer the interested reader to similar figures covering the rest of scenarios and the other indexes in the online resource accompanying this paper.

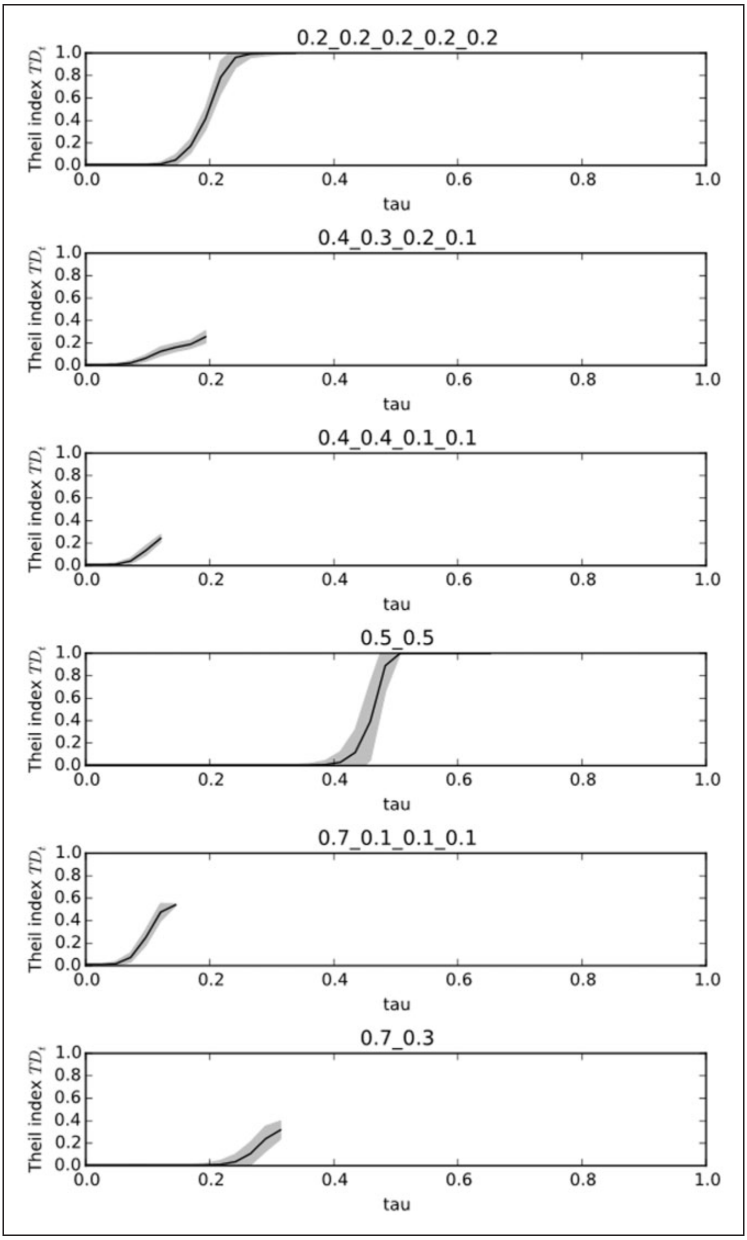

Figure 3. Theil Index behavior by population structure for the case of a "dense" geography with 25% vacancy rate (the line represents the average value across replications, while the shaded area covers plus/minus three standard deviations).

Visualizations such as those in Figure 3 are an excellent way to explore how the indexes considered behave under different scenarios. However, given the range of scenarios we consider in our simulations, they are not an effective tool for sensitivity analysis because that would involve comparisons across too many figures, plots, and lines. For a more systematic assessment, we turn to the second part of our results, where we summarize the simulation outputs in regressions that model the index obtained in a given successful draw as a function of the parameters used to obtain such a value. This method, generally known as response surface analysis (RSA; Box and Wilson, 1951), is implemented in the following way. For each of the four indexes, we calculate in each draw, we run the following regression

$$y_i = \alpha + \beta_1 \, VACR + \beta_2 \, SPRAWL + \beta_3 \, \tau + \beta_4 \, \tau^2 + \beta_5 \, \tau^3 + \sum_g \sum_k \delta_{gk} K_{gk} + \varepsilon_i \tag{1}$$

where y_i is the value of the index; α is the constant term; $VACR$ is a binary variable that takes one if the vacancy rate is 25%, zero if 15%; $SPRAWL$ is another binary variable taking one if the scenario is a ''sprawl'' configuration, zero if ''dense''; and τ is the value used to generate the map from which the value y_i was measured (which enters also with squared and cubic terms); K_{gk} are indicator variables that takes one if y_i is the value obtained for group g, under a scenario with population structure k, zero otherwise.

Results for the four regressions, one for each index, are displayed in Table 2. Given the very large number of observations, all coefficients are statistically significant at the 1% level. Overall, segregation index values tend to be higher for configurations with a 25% vacancy rate instead of 15%. This effect is significant but relatively small in magnitude and, based on exploration of the plots like that in Figures 2 and 3, are most likely due to the fact that higher vacancy rate allows not only for faster convergence but also for more complex scenarios to converge, which tend to have higher index values. ''Sprawl'' scenarios usually produce slightly higher index values for all cases except for the *EG* index. This effect is an order of magnitude greater than a varying vacancy rate.

The effect of the population structure is less straightforward to assess but can be discerned by considering the binary variables that mark the group and the structure. Since we are including an intercept and dropping the first group of the benchmark case, all the other coefficients should be interpreted as the effect of a particular case, *relative to* the benchmark situation. The first important aspect that stands out from the table is that in every case, every index value is significantly different from the benchmark case and associated with higher values on average. This is an important conclusion that can be extracted from the analysis: more complex population structures do indeed give rise to higher levels of clustering, ceteris paribus. Equally, scenarios with symmetric distributions of population yield symmetric results for each group. By examining further the magnitude of the coefficients within each column, it is possible to see that the highest values are attained in configurations with a larger number of groups. Additionally, within a given configuration, it is consistently the group with a smallest share of the population—the demographically smallest minority—that displays the highest value, and this effect is true no matter how many groups the configuration has.

Taking these patterns into consideration, it is no surprise that the one with highest values, for every index and group, is always the one we called ''diversity,'' where five different groups have the same share of population. A slightly different setup is necessary for the Theil index, which is a global measure that produces a single value for the entire map, instead of one per group. Such a feature requires us a slight modification of the regression, replacing the indicator variables for each group by one for the entire population. Despite this change, the same pattern arises: more diverse populations yield—for given preferences and geography—a greater degree of spatial clustering.

Inspired by the inspection of plots like Figure 3, we include τ in regression Equation (1) in a non-linear form. We introduce not only its value but also squared and cubic terms.

Table 2. Response surface analysis regression results.

	ellison_glaeser_egg_pop		isolation ii		segregation_gsg		theil _th	
	beta	sig	beta	sig	beta	sig	beta	sig
tau	0.182	***	5.0037	***	1.6917	***	−0.0626	*”*
taut	0.067	*A*	27.0463	***	1.543	***	6.4639	***
tau3	−0.3296	***	−40.2671	***	−1.6814	***	−5.1957	*”*
vacr	0.0041	***	0.0845	***	0.0201	***	0.0152	***
sprawl	−0.0194	**”	0.3193	***	0.107	***	0.0699	***
Intercept	−0.0294	***	−1.1476	***	−0.3343	***	−0.2503	***
gim 0.2_0. 2_0.2_0. 2_0.2lg0-0.20	0.0419	***	2.4817	***	0.4389	***		
gim 0.2_0. 2_0.2_0. 2_0.2lg1-0.20	0.0419	***	2.4828	***	0.4391	***		
gim 0.2_0. 2_0.2_0. 2_0.2lg2-0.20	0.0419	***	2.4829	***	0.4391			
gim 0.2_0. 2_0.2_0. 2_0.2lg3-0.20	0.0419	***	2.4841	***	0.4393			
gim 0.2_0. 2_0.2_0. 2_0.2lg4-0.20	0.0419		2.483	***	0.439	*-		
gim 0.4_0.3_0.2_0.1180-0.40	0.0263	***	1.5306	***	0.2658	***		
gim 0.4_0.3_0.2_0.1181 -0.30	0.0264	***	1.5539	***	0.2724	***		
gim 0.4_0.3_0.2_0.1182-0.20	0.0271	***	1.6561	***	0.314	***		
gim 0.4_0.3_0.2_0.1183-0.10	0.0397	***	2.4703	***	0.5272	***		
gim 0.4_0.4_0.1_0.11g0-0.40	0.0271	***	1.5941	***	0.2844	***		
gim 0.4_0.4_0.1_0.11g1-0.40	0.027	***	1.5937	***	0.284	***		
gim 0.4_0.4_0.1_0.11g2-0.10	0.039	***	2.3795	***	0.5198	***		
gim 0.4_0.4_0.1_0.11g3-0.10	0.039	***	2.378	***	0.5194	***		
gim 0.5_0.51g1-0.50	0		0		0			
gim 0.7_0.1_0.1_0.11g0-0.70	0.0266	***	1.59	***	0.3495	“**		
gim 0.7_0.1_0.1_0.11g1-0.10	0.0421	***	2.5206	**”	0.5398	***		
gim 0.7_0.1_0.1_0.11g2-0.1 0	0.042	***	2.5207	*-	0.5397	“*”		
gim 0.7_0.1_0.1_0.11g3-0.10	0.0422	***	2.5244	***	0.5405	“**		
gim 0.7_0.31g0-0.70	0.0117	***	0.8254	***	0.1572	***		
gim 0.7_0.31g1-0.30	0.0126	***	0.9119	***	0.1572	***		
pm 0.2_0.2_0.2_0.2_0.2							0.448	***
pm 0.4_0.3_0.2_0.1							0.2375	***
pm 0.4_0.4_0.1_0.1							0.2647	***
pm 0.7_0.1_0.1_0.1							0.3434	***
pm 0.7_0.3							0.1233	***
N	468,688		468,688		468,688		150,103	

*** is significant at the 1% level; ** at the 5%; and * at the 10%. “gim” stands for “group in mix” and relates to the group, within a given population structure, where an individual belongs. In the case of Theil, since it is a global measure, “pm” stands for “population mix,” and it relates to the overall structure set up.

To interpret the results in a visual way, we make use of Main Effect Plots, which display the direct effect on the dependent variable of modifying an independent variable, τ in this case (see Spielman et al., 2013, for a similar application of this tool). Figure 4 displays the plots for each of the four indexes analyzed. On the horizontal axis is the value of τ, while on the vertical one is that of the index. As expected, and already confirmed by Figure 3, there is a positive association between the two: as the level of co-location preference increases, the final degree of isolation/segregation/concentration/clustering increases also. However, the shape of this positive relationship varies substantially depending on the index. The *II* and *EG* indexes (first and third panels) display a fairly linear, albeit slightly concave, trend up until $\tau = 0.4$ and start to decrease slightly thereafter. This would suggest a peak after which the indexes, counterintuitively, start decreasing as segregation increases. However, it is important to note we have less certainty about this area of the value space as many draws are too complex to converge. The *SI* index (second panel) displays a simpler and more

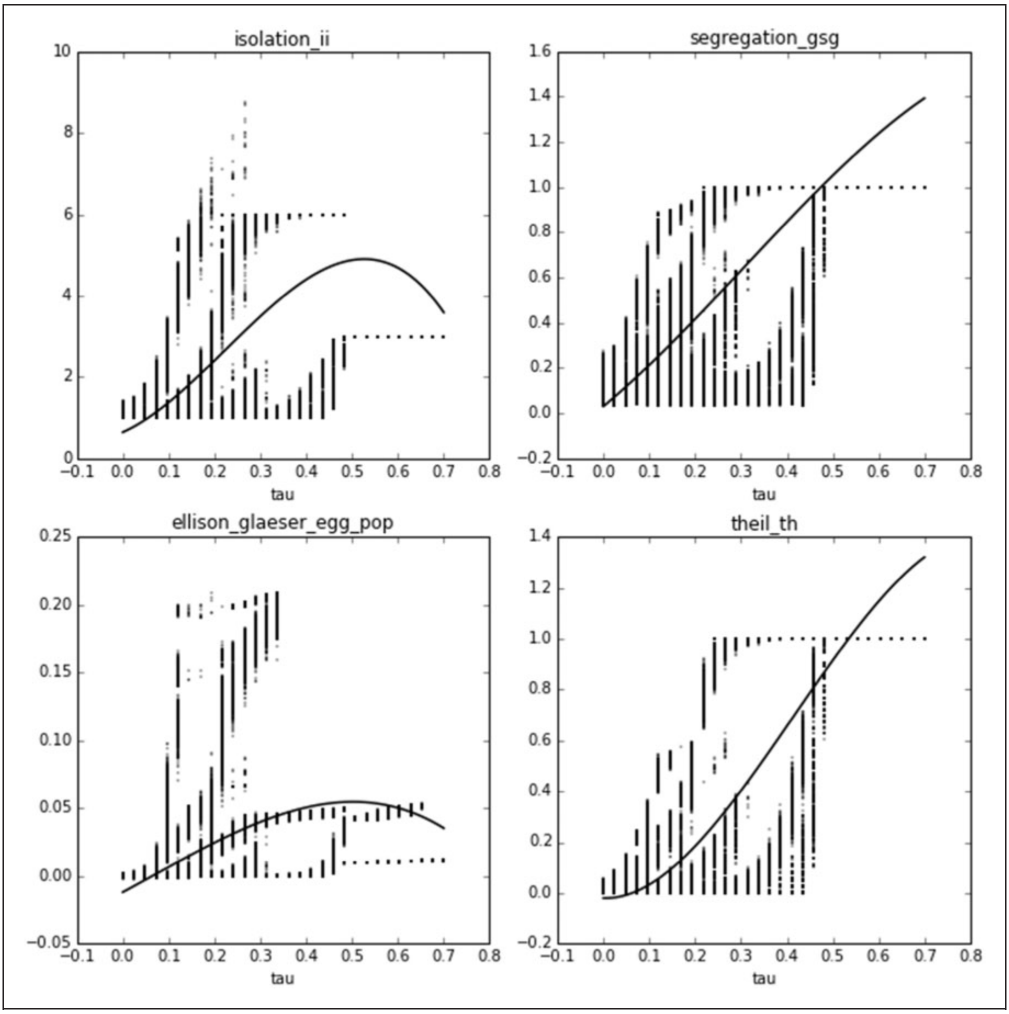

Figure 4. Main Effect Plots of τ on segregation indexes.

consistent positive relationship that is almost exactly linear. The *TD* index, displayed in the fourth panel, exhibits the *S*-shaped relationship already alluded to on the basis of Figure 3. Considering the four indexes altogether, both TD and SI appear more robust to higher degrees of segregation. SI, additionally, displays a more linear relationship when τ is lower, when diversity is high—a property that might be desirable in some contexts.

Concluding remarks

Cultural diversity has in recent decades become a fashionable term, but its measurement is still fraught with many problems. A major concern emerges from the fact that there is a great diversity in cultural diversity measures, either with respect to groups or areas (Nijkamp and Poot, 2015). In various cases, these measures have a normative framing, as their content is based on stylized or value-loaded interpretation of the composition of the relevant population. In a way, this is an issue comparable to the measurement of income distribution for which a range of indicators has also been developed. Clearly, the findings from a range of different cultural diversity measures will—by necessity and design—reflect potentially large differences in circumstances and interpretations, as was shown in our empirical work. We, therefore, tested whether our observed measures can be simply interpreted as revealing unobserved preferences for co-location with own type and found that the selected measures are indeed shown to be non-decreasing in increasing preference for within-group co-location. More diverse populations yield—for given preferences and neighbourhood structure—a greater degree of spatial clustering. We also conclude that the sensitivity of segregation measures to underlying conditions suggests that meaningful analysis of the impact of segregation requires spatial panel data modelling.

In the effort to keep the theoretical framework as simple as possible, we have made several simplifications. This is particularly true for the version of the Schelling model, we have used to generate maps and calculate the corresponding indexes. Sticking to the original barebones model keeps the setup tractable, intuitive and comparable to a large body of already established literature. However, some of the assumptions are arguably too simplistic and could be relaxed while still maintaining the general character of this approach. Some of the potential extensions include more sophisticated behavioral rules that account for several attributes such as labor markets, house prices, different generations and mobility costs. Additionally, an assumption of heterogeneity in behavior might imply that τ varies across groups or that it even can take negative values (i.e. a preference for heterophily rather than homophily).[11] Some of these situations have already been studied, but accounting for them when comparing the performance of segregation measures could yield new insights. Another complementary extension would go in the direction of blending the approach outlined here with microsimulation methods to explore its forecasting properties in real-world settings, with an eye toward a better understanding of the effectiveness of policy measures responding to segregation. A third set of extensions would pay attention to the dynamic *processes* that produce the outcome maps we study and the extent to which the evolution of observed segregation measures signal existence and attainability of equilibria or not. These are all very promising future extensions of this work.

In general, the conventional Schelling results on cultural segregation are confirmed with our simulation experiments. If we were to introduce time as an explicit dimension in the analysis, an even more complex situation is likely to arise. Although we have shed some light on the properties of some conventional segregation measures under a range of conditions, there continues to be a clear need for robust summary indicators of the spatio-temporal evolution of cultural diversity. Clearly, our findings have highlighted the need for a judicious use of indicators of the spatial distribution of specific groups of people in policy analysis and public debate.

Acknowledgements

Previous versions of this paper were presented at the 5th NORFACE Migration Conference, Radisson Blu Hotel, Berlin, Germany, November 2013; the 24th New Zealand Econometric Study Group meeting, Waikato Management School, University of Waikato, February 2014; and the PRSCO & SOCHER Pacific Conference of the RSAI and Sociedad Chilena de Economia Regional, Viña del Mar, Chile, August, 2015. Comments from Graeme Doole and other conference participants, and from three referees, are gratefully acknowledged.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Financial support was provided by the *NORFACE Research Programme on Migration*, www.norface-migration.org. Jacques Poot was also supported by the 2014–2020 New Zealand Ministry of Business Innovation and Employment (MBIE) grant UOWX1404 which funds the *Capturing the Diversity Dividend of Aotearoa New Zealand* (CaDDANZ) research programme, www.caddanz.org.nz

Notes

1. For a recent review and comparison of these methodologies, see e.g. Birkin and Wu (2012).
2. An earlier working paper version of this article reports the outcomes of the MC-ABM simulations in the context of the geographical distribution of the population across the city and wards of Amsterdam. In that case, we find that while segregation and isolation measures are still non-decreasing in the preference for co-location parameter, concentration measures are not. Hence given that the properties of some measures turn out to be sensitive to the assumed geography, in the present version, we adopt a more stylised geography and test the impact of the assumptions by means of response surface analysis.
3. See Nijkamp and Poot (2015) for a review of both kinds of perspectives.
4. In an earlier working paper version, we use the neighbourhoods of the municipality of Amsterdam as the geographical backdrop to the simulations. However, we could only capture population distribution in Amsterdam in a rather artificial and limited way, which may have impacted on the results. In the present paper, we use surface response analysis for a more systematic assessment of variation in the neighbourhood structure.
5. An investigation of the impact of the preference parameter varying across groups is outside the scope of the present paper but would be an interesting extension in future research.
6. Any ABM run is considered as non-convergent if after 2000 iterations no solution has yet been reached in which all of the agents are content. In practice, this implies that we are not making a distinction between cases in which there is not a feasible equilibrium and those where, even if there is an equilibrium, it is unlikely to be reached within a limited number of iterations.
7. It is simply due to the symmetry of the assumptions regarding the two groups in this specific simulation that the values of *SI*, *II* and *EG* are in Table 1 the same for both groups. The *TD* index is calculated for the total population and hence not group specific.
8. It also represents a computational challenge. The simulations were run on the Lisa system (https://www.surfsara.nl/systems/lisa). This is a computer cluster sponsored by a consortium of Dutch universities. We used 500 computing cores and still required a total running time of about 20 h. The computational burden also limited the range of sensitivity analyses that could be conducted. In the final section of this paper, we suggest several potential additional sensitivity analyses and extensions.
9. Similar plots for configurations that include a "sprawl" geography and 15% vacancy rate are not shown due to space constraints but are available as companion online resources to this paper. The same interpretation applies. See the companion website for this paper.
10. Given the other indexes often vary by reference group, plots for these focus either on the sample mean or standard deviation values, rather than condensing both into a single figure, as in the Theil case.
11. We thank a referee for suggesting these extensions.
12. A closely related segregation measure is the Gini coefficient GC_g, which measures the extent to which the group's population is differently distributed across areas, compared with the total

population (see e.g. Nijkamp and Poot, 2015). Consider a two-dimensional plot with cumulative population shares $P_{\bullet a}/P_{\bullet\bullet}$ of areas on the horizontal axis, ranked from the area which has the smallest share of g's population $P_{ga}/P_{g\bullet}$, and plot the cumulative shares of the group on the vertical axis. The curve that connects these cumulative shares of the total population and the cumulative shares of g's population is the well-known Lorenz curve and the area between the main diagonal and the Lorenz curve multiplied by two is the Gini coefficient. Given its similarity to the segregation index and space constraints, Gini coefficients of MC-ABM simulations are only reported in the online working paper version of this manuscript.

13. The modified segregation has been included in the clustering measures calculated for the simulations discussed in the online working paper.
14. Cutler et al. (1999), Maré et al. (2012) and Nijkamp and Poot (2015) suggest normalisations that bound the isolation index to be between zero and one. We do not apply such normalisations here.
15. A very similar index is the concentration index by Maurel and Sédillot (1999). This index has also been reported for the MC-ABM simulations discussed in the online working paper version but the index has been omitted here due to space constraints.
16. The Herfindahl index is a common measure of potential market power in an industry. It is equal to the sum of squared shares of firms in industry employment. The index ranges between one (in the case of monopoly) and approximately zero (in the case of many small firms).
17. This principle states that when a person moves from an area where his/her group share of the population is larger than in the area the person moves to, then the value of the segregation index declines. See Reardon and Firebaugh (2002) for further details.
18. This aspect of the model could be modified to include more sophisticated behavior, for instance by assuming the agents have macro information on the general pattern and use it to make more informed relocations. However, this is likely to result only in faster convergence to the stable equilibrium but not in qualitatively different outcomes. In order to keep the setup simple and comparable to most of the existing literature, we maintain random relocations.

References

Benenson I and Hatna E (2011) Minority–majority relations in the Schelling model of residential dynamics. *Geographical Analysis* 43(3): 287–305.

Benenson I, Hatna E and Or E (2009) From Schelling to spatially explicit modeling of urban ethnic and economic residential dynamics. *Sociological Methods & Research* 37(4): 463–497.

Benenson I, Omer I and Hatna E (2002) Entity-based modeling of urban residential dynamics: the case of Yaffo, Tel Aviv. *Environment and Planning B: Planning and Design* 29(4): 491–512.

Birkin M and Wu B (2012) A review of microsimulation and hybrid agent-based approaches. In: Heppenstall AJ, Crooks AT, See LM, et al. (eds) *Agent-Based Models of Geographical Systems*. London: Springer, p. 760.

Box GEP and Wilson KB (1951) On the experimental attainment of optimum conditions. *Journal of the Royal Statistical Society. Series B (Methodological)* 13(1): 1–45.

Brickell K and Datta A (2011) *Translocal Geographies: Spaces, Places, Connections*. Surrey: Ashgate Publishing.

Brown ZA and Moore RD (1970) The intra-urban migration process: a perspective. *Geografiska Annaler* B(52): 1–13.

Chen K, Irwin EG, Jayaprakash C, et al. (2005) The emergence of racial segregation in an agent-based model of residential location: The role of competing preferences. *Computational & Mathematical Organization Theory* 11(4): 333–338.

Clark WA (1991) Residential preferences and neighborhood racial segregation: A test of the Schelling segregation model. *Demography* 28(1): 1–19.

Clark WA (2006) Ethnic preferences and residential segregation. *Journal of Mathematical Sociology* 30(3): 319–326.

Clark WA and Fossett M (2008) Understanding the social context of the Schelling segregation model. *Proceedings of the National Academy of Sciences* 105(11): 4109–4114.

Crooks AT (2010) Constructing and implementing an agent-based model of residential segregation through vector GIS. *International Journal of Geographical Information Science* 24(5): 661–675.

Cutler DM, Glaeser EL and Vigdor JL (1999) The rise and decline of the American ghetto. *Journal of Political Economy* 107(3): 455–506.

Duncan OD and Duncan B (1955) Residential distribution and occupational stratification. *American Journal of Sociology* 60(5): 493–503.

Ellison G and Glaeser EL (1997) Geographic concentration in U.S. manufacturing industries: A dartboard approach. *Journal of Political Economy* 105(5): 889–927.

Feitosa FF, Le QB and Vlek PL (2011) Multi-agent simulator for urban segregation (MASUS): A tool to explore alternatives for promoting inclusive cities. *Computers, Environment and Urban Systems* 35(2): 104–115.

Fossett M and Dietrich DR (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.

Fossett M and Waren W (2005) Overlooked implications of ethnic preferences for residential segregation in agent-based models. *Urban Studies* 42(11): 1893–1917.

Kroeber AL and Kluckhohn C (1952) *Culture: A Critical Review of Concepts and Definitions*. Cambridge, MA: Peabody Museum.

Laurie AJ and Jaggi NK (2003) Role of 'vision' in neighbourhood racial segregation: A variant of the Schelling segregation model. *Urban Studies* 40(13): 2687–2704.

Maré DC, Pinkerton RM, Poot J, et al. (2012) Residential sorting across Auckland neighbourhoods. *New Zealand Population Review* 38: 25–54.

Massey DS and Denton NA (1988) The dimensions of residential segregation. *Social Forces* 67(2): 281–315.

Maurel F and Sédillot B (1999) A measure of the geographic concentration in French manufacturing industries. *Regional Science and Urban Economics* 29(5): 575–604.

Nijkamp P and Poot J (2015) Cultural diversity—A matter of measurement. In: Nijkamp P, Poot J and Bakens J (eds) *The Economics of Cultural Diversity*. Cheltenham: Edward Elgar, pp. 17–51.

Nijkamp P, Poot J and Sahin M (eds) (2012) *Migration Impact Assessment: New Horizons*. Cheltenham: Edward Elgar.

O'Sullivan D, MacGill J and Yu C (2003) Agent-based residential segregation: a hierarchically structured spatial model. In: *Proceedings of agent 2003 conference on challenges in social simulation*. Argonne National Laboratory and The University of Chicago, pp. 493–507.

Portugali J and Benenson I (1995) Artificial planning experience by means of a heuristic cell-space model. *Environment and Planning B: Planning and Design* 27: 1647–1665.

Reardon SF and Firebaugh G (2002) Measures of multigroup segregation. *Sociological Methodology* 32(1): 33–67.

Reardon SF and O'Sullivan D (2004) Measures of spatial segregation. *Sociological Methodology* 34: 121–132.

Samers M (2009) *Migration: Key Ideas in Geography*. London: Taylor & Francis.

Schelling TC (1969) Models of segregation. *The American Economic Review* 59(2): 488–493.

Schelling TC (1971) Dynamic models of segregation. *Journal of Mathematical Sociology* 1(2): 143–186.

Schelling TC (1978) *Micromotives and Macrobehaviour*. New York: Norton.

Simon JL (1999) *The Economic Consequences of Immigration*. Ann Arbor, MI: University of Michigan Press.

Spielman S and Harrison P (2014) The co-evolution of residential segregation and the built environment at the turn of the 20th Century: a Schelling model. *Transactions in GIS* 18(1): 25–45.

Spielman S, Yoo E-H and Linkletter C (2013) Neighborhood contexts, health, and behavior: understanding the role of scale and residential sorting. *Environment and Planning B: Planning and Design* 40(3): 489–506.

Theil H and Finezza AJ (1971) A note on the measurement of racial integration of schools by means of informational concepts. *Journal of Mathematical Sociology* 1: 187–194.

Urry J (2000) *Sociology Beyond Societies: Mobilities for the Twenty-first Century*. Routledge: International Library of Sociology.
van Mourik A, Poot J and Siegers JJ (1989) Trends in occupational segregation of women and men in New Zealand. *New Zealand Economic Papers* 23: 29–50.
Vertovec S (2007) Super-diversity and its implications. *Ethnic and Racial Studies* 30(6): 1024–1054.
Yin L (2009) The dynamics of residential segregation in Buffalo: an agent-based simulation. *Urban Studies* 46(13): 2749–2770.

Appendix 1

Measures of diversity

To define our measures, we need some notation. The number of people who belong to group g ($g = 1, 2,\ldots, G$) in area a ($a = 1, 2,\ldots, A$) is denoted by P_{ga}. The subscript $\bullet$ denotes aggregation over that index. Hence, $P_{\bullet\bullet}$ refers to the total population, with $P_{\bullet\bullet} = \sum_{g=1}^{G} \sum_{a=1}^{A} P_{ga}$.

The first measure we consider is the *dissimilarity index* introduced by Duncan and Duncan (1955). The index is a measure of displacement—the proportion of people in a group which would have to relocate in order to make their distribution identical to that of a reference group. When the dissimilarity index is computed between one group and all other groups combined, it is known as the *segregation index* (Nijkamp and Poot, 2015). The segregation index for group g across area units a is

$$SI_g = \frac{1}{2}\sum_{a=1}^{A}\left|\frac{P_{ga}}{P_g} - \frac{(P_a - P_{ga})}{(P - P_g)}\right| \qquad (2)$$

This index varies between 0 (the group's spatial distribution is the same as that of the rest of the population) and 1 (the group never co-locates with others in any particular area).[12]

One weakness of this index is that a particular value can arise with either a quite dispersed spatial distribution of the group or a rather clustered distribution, dependent on the spatial distribution of the population generally. There are other measures (such as the isolation index defined below) that inform on the extent to which the group of interest ''dominates'' the population of certain areas.

Another problem with SI_g is that the information it provides, i.e. the percentage of redistribution required to make a group distributed identical to the rest of the population, may have unrealistic implications for the distribution of population across areas that would result if the redistribution actually took place, particularly if the group under consideration is rather large. If we think of residential mobility being constrained by the available land in each area (which determines the maximum number of dwellings and hence the maximum population of the area), we could alternatively consider an index that measures the proportion of the population (both of the group of interest and of the rest of the population) that would need to shift to equalize the spatial distribution of the two groups, subject to the area populations remaining the same as before. Van Mourik et al. (1989) apply such a modified segregation index to the case of gender segregation across occupations. Because it can be shown that, for a given share of a group in the total population, the modified segregation index is a constant times the simple segregation index, we consider only the latter in the present paper.[13]

Because certain values of the segregation index can be obtained for both strongly or weakly clustered groups, we want to account for the exposure of the group to other people in the area. The *isolation index* captures the extent to which members of a population group are disproportionately located in the same areas, i.e. they are more clustered. Consider first the weighted average fraction (across all areas) of the population that belongs to group g, $\sum_{a=1}^{A} ga \frac{P_{ga}}{P_a}$. The weights are given by $_{ga} = \frac{P_{ga}}{P_g}$, i.e. the area shares of g's population (see e.g. Cutler et al., 1999). Clearly, $\sum_{a=1}^{A} ga = 1$ for each g. We now define the isolation index II_g simply as this average fraction divided by the group's share of the entire population

$$II_g = \sum_{a=1}^{A} \frac{ga \frac{P_{ga}}{P_a}}{\frac{P_g}{P}} \tag{3}$$

This measure captures the degree to which group members live in predominantly in areas in which they are over-represented. An isolation index value of 1 indicates that the group is distributed in proportion to the total population, while a much larger value of equal to P/P_a can be interpreted as total isolation (or extreme clustering) where all of the group locate in one particular area a in which there are no other groups.[14] Like the segregation index, the isolation index is a global measure that provides only limited information on the spatial patterns.

The next segregation measure that we consider is the Ellison and Glaeser (1997) *concentration* index denoted EG_g.[15] This index has been derived by means of a statistical model that compares the deviation of the *actual* correlation between location decisions made by pairs of members of the group (which can be positive or negative) with the correlation that would be observed when any individual is *randomly* assigned to any area (with the probability of being assigned to an area equal to the area's share of population P_a/P). This measure was originally derived to capture the geographic concentration of firms within an industry. We use a formula first applied by Maré et al. (2012), which differs slightly from the original formulation to reflect that the focus in our context is on people rather than firms. In the case of firms belonging to an industry (rather than people belonging to a cultural group), the difference in employment across firms is accounted for by a Herfindahl index.[16] Because all people in a group carry equal weight (unlike firms that are likely to differ in total employment), the Herfindahl index for a population group becomes $1/P_{g\bullet}$. A value of close to zero for EG_g indicates a lack of spatial concentration.

Using the same notation as before, the EG_g index is given by

$$EG_g = \frac{\frac{\left\{\sum_{a=1}^{A} (\frac{P_{ga}}{P_g} - \frac{P_a}{P})^2\right\}}{(1 - \sum_{a=1}^{A} (\frac{P_a}{P})^2)} - \frac{1}{P_g}}{(1 - \frac{1}{P_g})} \tag{4}$$

Ellison and Glaeser (1997) suggest that in order to determine a benchmark for their measure of concentration, the index should be first calculated for a group that could be considered to be roughly randomly allocated (proportional to area population sizes) across areas. In our present context of spatial segregation, the spatial distribution of females (or of males) relative to the total population may be used as a benchmark for comparison. Hence, for both genders EG_g will have low values which can be used as a benchmark for gauging spatial segregation of culturally defined groups.

The measures of spatial clustering defined above have all been designed to measure the spatial clustering of a specific population group vis-à-vis the rest of the population. Such measures are particularly useful, and unambiguous, when the population of interest is considered dichotomous: migrant versus non-migrant, indigenous versus non-indigenous, black versus white, etc. Given the growing diversity of population across range of types with any given domain (ethnicity, religion, etc.), there is a growing need for multigroup summary measures. Such measures have been reviewed by Reardon and Firebaugh (2002) on the basis of seven criteria. They conclude that the information theory-based index proposed by Theil and Finezza (1971) is the most satisfactory index in terms of their criteria. This index has two particularly desirable properties. First, it is the only one that satisfies the principle of transfers.[17] Second, it is the only multigroup index that can be decomposed into a sum of between- and within-group components. Hence, we include Theil's information index of diversity TD in our range of clustering measures. This index can be calculated as follows

$$TD = \sum_{g=1}^{G}\sum_{a=1}^{A}\frac{P_a}{P}\left[\frac{\frac{P_{ga}}{P_a}\left(\ln\left(P_g/P\right) - \ln\left(P_{ga}/P_a\right)\right)}{\sum_{g=1}^{G}(P_g/P)\ln(P_g/P)}\right] \tag{5}$$

The TD index values vary between zero, when all groups are equally distributed across areas and approximately one, when there is near-complete segregation (note that the index calculation requires there to be at least one person of each group in each area). Computationally, since the shares cannot take zero values, we add a negligible amount to zero counts before calculating the index.

Appendix 2

The Schelling model

The Schelling (1969, 1971, 1978) model is usually credited as one of the first ABMs and had a tremendous influence on the subsequent literature, its effect spanning even up to today. The basic intuition of the model is rather straightforward. A given number of agents are assigned to different groups (e.g. natives and migrants) and randomly spread across a residential space divided up into several areas or (bounded) neighborhoods. Every individual follows a simple rule of behavior: given the preferences set, if the proportion of individuals in her neighborhood is within her preferences, she is content and thus stays put; if it is not, the individual will decide to relocate once the model starts. The evolution of the model is sequential: in a series of steps, discontented agents are randomly assigned a new location and their status is evaluated again.[18] Once every single agent is satisfied with the composition of the neighbourhod, the model has converged and a solution has been found. As suggested in Schelling (1971) and re-illustrated in Clark (1991), the existence of equilibrium in a model with a particular set of parameters can be understood as the overlap of compatible outcomes for each population group in the model. This can be thought of as a portion of the space of all the possible outcomes of the model in which each group would be satisfied. In this context, it is possible to conceive some situations in which such space is an empty set and there is no possible solution that pleases all the agents. In this case, the model would never converge. It is also easy to see that the chance that a model does not converge increases as the complexity in the setup (number of groups, combination of proportions, tolerance rate, etc.) becomes greater.

The Schelling model has had a tremendous impact on the ABM literature, particularly on that focused on residential segregation issues, including even relatively recent studies (see Clark, 2006). Although providing a comprehensive overview of the literature is beyond the scope of the present paper, let us mention some of the most relevant articles for the present study. Three strands of literature can be distinguished. One set of articles has tried to test the validity and robustness of the model's predictions in comparison to real world data. Two good references of this approach are Clark (1991), who delves into some of the theoretical implications and their relation with the empirical evidence, and Clark and Fossett (2008), who present an overview of the main contributions from the social sciences in terms of gaining understanding about the behavioral forces underlying the Schelling model.

Another branch of the literature has focused on extending the original model. Among these contributions, O'Sullivan et al. (2003) present a variation combining two different definitions of neighborhoods. Laurie and Jaggi (2003) propose an extension that modifies the ''vision'' parameter and find that results change substantially, although this finding is contested by Fossett and Waren (2005). Portugali and Benenson (1995) analyse the group-specific movement thresholds and the residential migration symmetry assumption. Chen et al. (2005) include income, neighborhood attributes and asymmetric preferences. Fossett and Dietrich (2009) show that the results of the Schelling model are in many ways robust to the effects of city size, shape and form, and neighborhood size and shape. Crooks (2010) proposes the introduction of elements of real geography (fences, boundaries, irregularities, etc.) through the combination of agent-based approaches with GIS. Benenson and Hatna (2011) explore the properties when the number of groups and neighborhood sizes change, focusing particularly on the model's dynamic aspect.

Finally, a broader body of work applies the main mechanics of the model to particular problems. Benenson et al. (2002, 2009) use the model to study residential dynamics in an Israeli neighborhood and test whether its predictions reflect closely the actual outcomes, as made available by detailed geo-referenced micro-data. They find very robust behavior of the model. Feitosa et al. (2011) present a simulator that allows them to explore different extensions of the original setup and apply it to a case study of a Brazilian city. Yin (2009) couples a basic model that includes house prices with real data from Buffalo (NY) to replicate the actual segregation patterns of that city; while Spielman and Harrison (2014) also include income and replicate the conditions of Newark in the year 1880 to consider the effect of changes in urban spatial structure. The general result in relation to the validity of the model in real world contexts is that, albeit in a simplistic view, the Schelling model effectively captures several elements of human behavior.