How Diverse can Spatial Measures of Cultural Diversity be? Results from Monte Carlo Simulations on an Agent-Based Model

Daniel Arribas-Bel[1]
Peter Nijkamp[2]
Jacques Poot[3]

[1] University of Birmingham, United Kingdom;
[2] Faculty of Economics and Business Administration, VU University Amsterdam, Tinbergen Institute, the Netherlands;
[3] The University of Waikato, New Zealand.

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

Daniel Arribas-Bel [a,†] , Peter Nijkamp [a]   and Jacques Poot [b]

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 – are often only capable of identifying just one aspect of this distribution. The strengths or weaknesses of any measure can only be comprehensively assessed empirically. This paper provides evidence on the empirical properties of various spatial measures of cultural diversity by using Monte Carlo replications of agent-based modeling (MC-ABM) simulations with synthetic data assigned to a realistic and detailed geographical context of the city of Amsterdam. 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 number and shares of various cultural groups, and their preferences with respect to co-location. Our MC-ABM data generating process generates output maps that enable us to assess the performance of various spatial measures of cultural diversity under a range of demographic compositions and preferences. We find that, as our simulated city becomes more diverse, stable residential location equilibria are only possible when particularly minorities become more tolerant. We test whether observed measures can be interpreted as revealing unobserved preferences for co-location of individuals with their own group and find that the segregation and isolation measures of spatial diversity are shown to be non-decreasing in increasing preference for within-group co-location, but the Gini coefficient and concentration measures are not.

Keywords: cultural diversity, spatial segregation, agent-based model, Monte Carlo simulation

JEL-classification: C63, J15, R23, Z13

Version: April 17, 2014

[†] Corresponding author. School of Geography, Earth and Environmental Sciences, University of Birmingham (UK). Email: D.Arribas-Bel@bham.ac.uk
[a] Department of Spatial Economics, VU University Amsterdam, Amsterdam, The Netherlands
[b] National Institute of Demographic and Economic Analysis, University of Waikato, New Zealand.

1. Introduction: cultural diversity and space-time dynamics

International migration is not predominantly a demographic phenomenon of population movement from one country to another. It 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 host 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, or other migrant groups, or local natives. Consequently, the social geography of a city may exhibit drastic changes in its socio-cultural composition and the spatial distribution of culturally-defined groups over several decades after an initial migration influx. Cultural diversity triggers a whole range of spatial socio-economic phenomena related to housing and labour 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 analysed by means of the so-called Schelling model (Schelling, 1969; 1971; 1978).

The background of Schelling's compositional dynamics model is formed by the anticipated utility derived from population composition in urban neighbourhoods. Cultural diversity is – beyond a critical threshold – a key factor in spatial movement. 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 behaviours of residents in a city, whose location is often interpreted as the outcome of a process of minimising or resisting so-called residential stress (Brown and Moore, 1970). If this cultural diversity is related to a wealth of differences across nationalities (especially in modern mega-cities such as London or New York), one may even speak of super-diversity (Vertovec, 2007).

The 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 will only marginally change, but this move may prompt a chain reaction of other people in the same neighbourhood. Such a chain phenomenon is a clear example of social contagion. The simulation of such dynamic change patterns 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 interactive studies on dynamic human behaviour, often including the question whether a new social and spatial equilibrium may emerge after a small initial disturbance.

In this paper we are concerned with measuring the spatial diversity of the cultural identity of the population, whether measured subjectively or objectively. We consider various common measures of spatial diversity in cultural identity. However, such measures – such as segregation, isolation or

---

[1] For a recent review and comparison of these methodologies, see e.g. Birkin and Wu (2012).

concentration indexes – are often only capable of identifying just one aspect of the spatial distribution of the population across culturally-defined groups. The strength or weakness of any spatial diversity measure can only be comprehensively assessed empirically. For this we use the classical agent-based model (ABM) from Schelling (1971) as the "theoretical engine" to generate several patterns of spatial segregation. The data inputs include the initial population, the number and shares of various cultural groups, and their preferences with respect to co-location. We embed these data in a realistic and detailed geographical context: the city of Amsterdam. We provide evidence on the empirical properties of various cultural diversity measures by using Monte Carlo (MC) replications of each ABM simulation, yielding what we will refer to as the MC-ABM approach. Our MC-ABM data generating process generates output maps that enable us to assess the performance of various spatial segregation measures under a range of demographic compositions and preferences. The combination of agent-based modelling with Monte Carlo simulations – two usually separately employed techniques – and the application of MC-ABM to the evaluation of diversity measures represents the main contribution of this paper.

We find, not unexpectedly, that for a given geographical configuration of neighbourhoods certain population distributions and cultural preferences with respect to co-location are more compatible with equilibrium in residential choices than others. As societies become more diverse, stable residential location equilibria are only possible when people become more tolerant. Specifically, there appears to be a "rule of the minority" in that the minority "controls" how intolerant preferences can become for equilibrium to be feasible. More importantly, given our focus on the interpretive value of various spatial measures of cultural diversity, we test whether our observed measures can be interpreted as revealing unobserved preferences for co-location with own type. We find that the segregation and isolation measures of spatial diversity are shown to be non-decreasing in increasing preference for within-group co-location, but the Gini coefficient and concentration measures are not.

This paper is organized as follows. The next section is devoted to a concise review of our selected spatial indicators of cultural diversity that form the foundation for our empirical work. In Section 3 the principles of the methodology employed in this study are highlighted, with a particular emphasis on the classical Schelling model. Next, Section 4 offers a concise description of the simulation design and strategy which produces the main results. These results are extensively presented in Section 5. Section 6 concludes.

2. Measuring cultural diversity across space

Culture is a complex socio-anthropological phenomenon that concerns certain common behaviour, 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. 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. This kind of cultural diversity is also referred to as spatial segregation, dissimilarity or clustering. It is commonly measured in terms of residential location, but can also be measured in terms of workplace location or other location-referenced activities.

There are many ways to describe spatial clustering (e.g. Massey and Denton, 1988; Reardon and O'Sullivan, 2004). Measures can be *local* (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* (the extent to which a group is distributed differently across areas as compared with the population overall). They can also explicitly take the *location* and *proximity* of various areas into account. In this case there can be again local measures of dissimilarity (e.g., the extent to which the share of a group in a particular area differs from that of neighbouring areas) or global measures (e.g., the average correlation between a group's share in an area and the group's share in surrounding areas). Here we are concerned with global summary measures that describe the dissimilarity of the distribution of groups across predefined areas, but without taking distance-defined spatial correlation into account.

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

The first global spatial diversity 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 one 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*. The segregation index for group g across area units a at time t is

$$SI_{gt} = \frac{1}{2} \sum_{a=1}^{A} \left| \frac{P_{gat}}{P_{g\bullet t}} - \frac{(P_{\bullet at} - P_{gat})}{(P_{\bullet\bullet t} - P_{g\bullet t})} \right| \tag{1}$$

This index can of course also be used to measure segregation across another classification, such as occupation. The 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).

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_{gt} 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 for each area (which determines the maximum number of dwellings and hence the maximum population in the area), it is preferable to 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 equalise the spatial distribution of the two groups, subject to the area populations remaining the same as before. Van Mourik et al. (1989) apply such an index to the case of gender segregation across occupations. Clearly, the same issue arises there as the redistribution of the workforce across occupations to ensure gender equality should not be such that the total number of jobs in each occupation (which is determined by labour demand) changes.

In our spatial segregation context, the modified segregation index is given by

$$MS_{gt} = \frac{1}{2P_{\bullet\bullet t}} \sum_{a=1}^{A} \left(\left| P_{gat} - P_{g\bullet t} \frac{P_{\bullet at}}{P_{\bullet\bullet t}} \right| + \left| (P_{\bullet at} - P_{gat}) - (P_{\bullet\bullet t} - P_{g\bullet t}) \frac{P_{\bullet at}}{P_{\bullet\bullet t}} \right| \right) \tag{2}$$

Van Mourik et al. (1989) show that the modified segregation index MS_{gt} is equal to the conventional segregation index times a factor that depends on the fraction of the total population that is in group g:

$$MS_{gt} = 2 \frac{P_{g\bullet t}}{P_{\bullet\bullet t}} \left(1 - \frac{P_{g\bullet t}}{P_{\bullet\bullet t}} \right) SI_{gt} \tag{3}$$

This index is always smaller than SI_{gt} and reaches a maximum value of 0.5 when group g accounts for half the population and is completely segregated from the rest of the population.

Because, as noted above, certain values of the (modified) segregation index can be obtained whether the group is strongly clustered or not, we want to account for the exposure of the group to other people in the area. Our third spatial segregation index is the *isolation index*. This 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, $\sum_{a=1}^{A} \omega_{gat} \frac{P_{gat}}{P_{\bullet at}}$. The weights are given by $\omega_{gat} = \frac{P_{gat}}{P_{g\bullet t}}$, i.e. the area shares of g's population (see e.g. Cutler et al., 1999). Clearly, $\sum_{a=1}^{A} \omega_{gat} = 1$ for each g and all t. We now define the isolation index II_{gt} simply as this average fraction divided by the group's share of the entire population:

$$II_{gt} = \sum_{a=1}^{A} \omega_{gat} \frac{P_{gat}}{P_{\bullet at}} \bigg/ \frac{P_{g\bullet t}}{P_{\bullet\bullet t}} \tag{4}$$

This measure captures the degree to which group members live 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 $P_{\bullet\bullet t}/P_{\bullet at}$ 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.[2] Like the segregation index, the isolation index is a global measure that provides only limited information on the spatial patterns.

---

[2] Cutler et al. (1999) and Maré et al. (2012) suggest normalisations that bound the isolation index to be between zero and one. We do not apply such normalisations here.

The fourth segregation measure we consider is the Gini coefficient GC_{gt}, which measures the extent to which the group's population is differently distributed across areas as compared with the total population. This index is therefore quite similar to the segregation index (see e.g. Massey and Denton, 1988). Consider a two-dimensional plot with cumulative population shares $P_{\bullet at}/P_{\bullet \bullet t}$ of areas on the horizontal axis, ranked from the area which has the smallest share of g's population $P_{gat}/P_{g\bullet t}$, 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 divided by two is the Gini coefficient.

The final two segregation measures that we consider are the closely-related Ellison & Glaeser (1997) and Maurel & Sédillot (1999) *concentration* indexes, denoted EG_{gt} and MS_{gt} respectively. Both formulae have 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 an area being assigned equal to the area's share of population $P_{\bullet at}/P_{\bullet \bullet t}$). The measures were originally derived to capture the geographic concentration of firms within an industry. We use formulae first applied by Maré et al. (2012), which differ slightly from the original formulations 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.[3] Because all people in the group carry equal weight (unlike firms that are likely to differ in employment), the Herfindahl index for a population group becomes $1/P_{g\bullet t}$ (as in the unweighted index of Maurel & Sédillot, 1999). A value of zero for either of these indices indicates a lack of spatial concentration. The two indices differ only slightly in value, but they have different properties (see below).

Using the same notation as before, the EG_{gt} index is given by

$$EG_{gt} = \frac{\frac{\left\{\sum_{a=1}^{A}\left(\frac{P_{gat}}{P_{g\bullet t}} - \frac{P_{\bullet at}}{P_{\bullet \bullet t}}\right)^2\right\}}{\left(1 - \sum_{a=1}^{A}\left(\frac{P_{\bullet at}}{P_{\bullet \bullet t}}\right)^2\right)} - \frac{1}{P_{g\bullet t}}}{\left(1 - \frac{1}{P_{g\bullet t}}\right)} \tag{5}$$

while the MS_{gt} index is given by

$$MS_{gt} = \frac{\frac{\left\{\sum_{a=1}^{A}\left(\frac{P_{gat}}{P_{g\bullet t}}\right)^2 - \sum_{a=1}^{A}\left(\frac{P_{\bullet at}}{P_{\bullet \bullet t}}\right)^2\right\}}{\left(1 - \sum_{a=1}^{A}\left(\frac{P_{\bullet at}}{P_{\bullet \bullet t}}\right)^2\right)} - \frac{1}{P_{g\bullet t}}}{\left(1 - \frac{1}{P_{g\bullet t}}\right)} \tag{6}$$

Ellison & Glaeser (1997) suggest that in order to determine a benchmark for their measure of concentration, the index could be 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

---

[3] The Herfindahl index is a common measure of potential market power in an industry that 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).

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_{gt} (and MS_{gt}) will have low values which can be used as a benchmark for gauging spatial segregation of culturally defined groups.

Maré (2005) suggests that an intuitive interpretation of the MS_{gt} index is the extent to which the distribution of the various area shares of group g ($P_{gat}/P_{g \bullet t}$), is steeper than that of the distribution of the total population across areas ($P_{\bullet at}/P_{\bullet \bullet t}$). If so, MS_{gt} will be positive. If the distribution of group g is flatter than that of the total population, MS_{gt} will be negative. In the case of the EG_{gt} index, both steeper and flatter deviations from the distribution of the total population increase the value of the index.

3. The Schelling model as the theoretical engine

In order to examine the behaviour of the various indicators of spatial diversity just reviewed under a range of assumptions, our approach relies on computer simulations. We evaluate the indicators calculated 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 spatial distribution maps is the original segregation model of Schelling (1969, 1971, 1978), credited as one of the first Agent-Based Models (ABMs), and the first one to tackle spatial patterns of segregation. In particular, we adopt the so-called "bounded neighbourhood" variation. This is both 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.

The basic intuition of the original model is rather straightforward. A handful of agents assigned to different groups (e.g. blue and red) are randomly spread across a geography divided up into several areas or (bounded) neighbourhoods. Every individual follows a simple rule of behaviour: given the preferences set, if there are enough individuals of his or her type in the neighbourhood, (s)he is content and thus stays put; if there are not, the agent 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[4]. Once every single agent is satisfied with the composition of the neighbourhood, 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

---

[4] 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 would likely 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.

the chance of a model not being able to 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, spanning even to 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 groups 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 behavioural 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 neighbourhoods. 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, neighbourhood attributes and asymmetric preferences. 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 neighbourhood sizes change, focusing particularly on its dynamic aspect.

Finally, a broader body of work takes the main mechanics of the model and applies these to particular problems. Benenson et al. (2002, 2009) use the model to study residential dynamics in an Israeli neighbourhood and test whether its predictions reflect closely the actual outcomes, as made available by detailed geo-referenced micro-data, finding very robust behaviour of the model. Feitosa et al. (2011) present a simulator that allows them to explore different extensions of the original setup proposed by the literature and apply it to a case-study of a Brazilian city. Yin (2009) couples a basic model that also 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.

A few parameters in the model are key in determining the outcome. These include: the spatial layout adopted, the size and composition of the population and the individual residential preferences. In our experiments, we keep the first and second fixed while modifying the composition of the population and its preferences to obtain urban distributional maps with different characteristics.

The choice of geography is among the first to be made when designing a Schelling-type model. Boundaries that create the neighbourhoods in which agents will locate and will interact with others must be defined. Although most examples in the literature employ pixel grids, we follow a recent suggestion by Crooks (2010) and introduce elements of a real geography through vector GIS. In particular, we use the official neighbourhoods (so-called "buurten") of the municipality of

Amsterdam, which are shown in Figure 1 overlaid on a map of the city. This provides a realistic geographical backdrop to our theoretical parameters and simulations. The set of neighbourhoods is made up of 96 polygons which form the potential locations for the agents in the models. The number of potential sites is set to 4,900, and they are distributed across the polygons based on the area of the polygons. This is reasonable because, for constant density, larger areas will have more dwellings. Because there must be opportunities for people to move, not all sites are occupied. In accordance with many other examples in the literature, the initial distribution allows for 25 percent of sites being vacant. This ensures that agents can move more easily around and find locations where they are content to stay. Hence, equilibrium is achieved faster, as long as it is feasible.

Figure 1 about here

It is useful to point out that the municipality of Amsterdam (which is part of the larger Amsterdam agglomeration) covers an area of about 219 km^2. Consequently, each of our polygons covers about 2.28 km^2 on average and each potential site for agents is on average about 45,000 m^2. The population of the city of Amsterdam is about 810,000 people. Hence the population density is 3,699 persons per km^2, which is equivalent to 270 m^2 per person. Even with our advanced computing capabilities we are unable to simulate the behaviour of *all* of Amsterdam's 810,000 residents. Instead we consider 3,675 synthetic agents who, in the context of the actual population of Amsterdam, would each represent a group of 220 culturally homogeneous residents occupying a site of 45,000 m2, i.e. each resident occupies an average area of about 200 m^2, which is – of course by construction – 75 percent of the gross area, given that 25 percent of sites are vacant.

The cultural composition of the population of agents is defined by two parameters: the number of groups and the shares they represent. These two are key instruments to determine the outcome patterns that we need to evaluate the properties of several spatial diversity indices. 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 neighbourhood 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 former implies the agent is always content and the latter implies that he or she is satisfied only when everyone else in the neighbourhood belongs to the same group.

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. In addition, higher requirements imply it is harder for the model to converge into a solution in which every agent is content. 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 neighbours to be alike in order not to leave the neighbourhood. 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 the particular scenario is considered finished.[5]

---

[5] Any ABM run is considered as non-convergent if after 2,000 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

Once the model has been run with a set of parameters (i.e. a combination of the number and shares of cultural groups and a value for τ), it is possible to visualize the outcomes on a map. Figure 2 displays two outcomes of the model in the same scenario but with very different preferences. Both maps represent a setting with only two groups, blue and red, who are assigned each half of the population. On the left side, individuals do not have a preference for the characteristics of their neighbours and, consequently, any outcome (even the initial random one) represents an equilibrium distribution of the 3,675 agents (the 1225 vacant sites are also clearly visible). The pattern in this case is that of a random spread with colours not showing any clear spatial clustering. The right panel displays a similar setup but, in this case, individuals have a strong preference for similar neighbours. In particular, at least 65 percent of the neighbourhood population needs to be of the same group for an individual to be content. This translates in a clearly distinct spatial pattern with a high degree of co-location and spatial concentration. Red agents are surrounded by other red peers, while blue ones locate among other blue individuals. With a 50/50 split of the population, a preference for 65 percent of the neighbourhood being of one's own type may be expected to be difficult to achieve but the assumption of one quarter of sites being vacant is sufficient for the majority of simulations to yield equilibria (including the one displayed in Figure 2).

Figure 2 about here

Although illustrative, the outcome maps of Figure 2 are not the main end of this paper. Our ultimate goal is to study different measurements 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 underlying them are used to calculate the six diversity indices presented in Section 2. Continuing the illustration, Table 1 shows the measures calculated for the two maps displayed in Figure 2. These have been computed for each group (red and blue) using the counts of agents across each of the 96 neighbourhoods and 4900 sites, resulting from the Schelling processes. It can be clearly seen that all indices reflect 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.[6] This is the *output* that we analyse and base our conclusions on. Before that, we need to create a robust sample of such results that allows us to draw sensible conclusions. For that reason, we turn to a Monte-Carlo approach, the other key computational tool making possible in this study.

Table 1 about here

4. A Monte-Carlo agent-based approach

In order to compare the properties of different indices of spatial segregation in a synthetic but empirical context, a handful of solutions to the Schelling model are not sufficient. We need a larger,

---

in which there is not a feasible equilibrium and those where, even if there is, this is highly unlikely and thus is not reached in the several trials we attempt before considering finished and moving on to the next setting.
[6] This is harder to see for the Ellison and Glaeser index EG_{gt} and the Maurel & Sedillot index MS_{gt}, which are very small numbers. However, the numbers of the right hand side are in absolute value three to seven times those on the left hand side.

more robust, solution 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 extremely unlikely initial conditions may have on the conclusions; and, two, to cover an extensive range of cases and situations that lead us to general conclusions. The tool that allows us to meet these objectives is a Monte Carlo approach. This simulation method, combined with the Agent-Based model used at the core of this strategy, yields what we referred to in the introduction as the Monte-Carlo Agent-Based Modelling (MC-ABM) methodology.

The idea is simple: to conceive a set of scenarios by varying the number and proportions of groups, and replicate them over a wide spectrum of τ values. Each of these combinations are then replicated a sufficient number of times to stable results. For each random draw, the spatial segregation indices discussed in section 2 are calculated and stored. Average values across replications are then represented in an organized way that permits direct visual comparison, not only within but also across different scenarios. This highlights the main features on which we base conclusions and recommendations.

Every one of the simulated scenarios is defined by a combination of three parameters: preferences of agents, number of groups, and share of the population in each of these groups. For every triad, we run the model 500 times and obtain a set of indices for each run. This produces a substantial amount of output, which is very hard to explore and understand in raw form.[7] We adopt a visualization approach that allows for immediate comparison and for meaningful exploration of the performance of the various segregation measures across scenarios. Specifically, for every population structure, we generate a figure that contains one plot per indicator. Each of these plots displays different values of τ (preferences) along the X axis, and the *average* value for each group of the diversity measure obtained in the up to 500 replications in that τ-population setting on the Y axis (the average is calculated only for cases which yielded equilibria).

5   Results

When designing the scenarios to simulate, it is important to balance the trade-off between obtaining a wide and general range of situations and population characteristics and keeping the number of cases small enough so the computational burden remains realistic. With this goal in mind, we carefully select six combinations of number and share of groups in which the population is allocated. These are the following: a benchmark of two groups with equal shares (50%-50%); two groups with 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%), to which we will refer 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 respectively diminishing presence (40%-30%-20%-10%), referred to as the "ladder"; and the "diversity" case with five groups of equal weight (20% each).

---

[7] It also represents a computational challenge. The simulations were run on the Lisa system (https: //www.surfsara.nl/systems/lisa), a compute cluster sponsored by a consortium of Dutch universities. We used 500 computing cores and still required a total running time of about 20 hours.

In all of these cases, we simulate 500 independent maps for every value of τ we test, which begins always with 0, increases by 0.024 and stops when the model does not converge (the maximum we run is $\tau = 0.7$). As mentioned above, the chances of a model converging are related to the feasibility of equilibrium for a particular case. On one side of the spectrum, when $\tau = 0$, every 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 neighbours are required to be alike for an individual to be content, there could be cases in which the 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 the model needs to go through in order to find a solution, if it does at all. Figure 3 displays, for each scenario, the number of successful runs and the average number of iterations that the model needed to complete (topping at 2,000, which is the limit for convergence after which, for computational reasons, we skip the scenario). As can be quickly seen, the more groups there are and the smaller a minority is, the more difficult it is to find a solution and, hence, the less restrictive preferences can become. 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. Although simple and obvious in reality, it is comforting to find a clear counter-part in our MC-ABM simulations arising from such a simple set of assumptions.

Figure 3 about here

As mentioned above, the amount of information generated by simulations represents a challenge to organize and visualize. We solve it by structuring it in a set of charts arranged in a familiar way that allows for meaningful comparison. Figures 4-9 contain results for each of the scenarios described above. This is meant to replicate the kind of situation a researcher would face when analysing real data: one single entity (e.g. a region, a city) is studied and several indices are calculated on the same data. The figure is composed of six subplots, one for every index described in Section 2. These graphs in turn display information for the average value of the index over the at most 500 replications run on each value of τ, ascending along the horizontal axis. This layout permits to interpret the plot as the evolution of a given index, for each group in a given scenario, as the setting becomes more and more restrictive and the population requires a higher proportion of neighbours alike to be content in a given location.

As an illustration of how to consider and interpret these figures, we can use the benchmark case. It is important to remember each subfigure is based on the exact same set of data, so differences in the shape and evolution of the graphs can be attributed to differences in the index they reflect. Equally, it is important to note that, although the vertical axis is expressed on a different scale for each index, the horizontal one is standardized across plots of the same figure, allowing us to compare directly across plots the behaviour of indices as τ increases. Between figures, we change the scale to obtain a better view of the differences between indices in the τ-space in which the model converges for that particular scenario. As it can be seen from Figure 4, different indices portray fairly different views from the same maps. The index values are the same for the two groups, because they both represent half the population, but how they change as the agents' preferences become more restrictive is particular to the index. Both segregation measures (SI_{gt} and MS_{gt}) and the isolation index II_{gt} offer a similar view: values stay very low when agents are content with almost any outcome but, once they require about 35% of the neighbours to be of the same group, the indices

sharply increase until reaching their top value (1, 0.5 and 2, respectively) when half the neighbours are required to be of the same group, after which the indices remain at their maximum values (reflecting total segregation of the two groups: neighbourhoods are either entirely of one type or of the other). The Gini coefficient offers a slightly different perspective: its increase is more progressive and with a smoother slope, closer to a linear relation, up to its peak at almost 0.8, only to slightly decrease afterwards and remain stable at around 0.72. This effect is much more pronounced in the case of the Ellison & Glaeser (EG_{gt}) and Maurel & Sedillot (MS_{gt}), both with a very similar trajectory: they remain mostly unaffected by the increase in τ up to a threshold around 0.43, where they experience a sharp rise and decrease almost completely to the initial levels when τ reaches 0.7. Essentially this is because, with complete segregation due to high values of τ, half of the neighbourhoods are entirely occupied by one group and half by the other. With a large number of neighbourhoods, the spatial distribution of the total population then becomes equally "steep" or "flat" as the spatial distribution of either of the groups.

Rather than replicating a similar description for every single scenario in all the other cases, we focus in the remainder of this section on extracting stylized facts about the behaviour of these six indicators that can teach us important lessons for their application in real world contexts.

A first general pattern that appears is that scenarios with symmetric distributions of population yield symmetric results for each group in the map, with the exception of the concentration measures, for which this is true only in approximate terms. Equally, asymmetric settings do not allow for more restrictive preferences. In a way, this result can be understood along the same lines as the chance of convergence for the model before: not only more diverse, but also more complex social configurations (e.g. situations with a minority and a majority population) require larger degrees of tolerance in their members for a stable equilibrium to be reached. Again, this is fairly intuitive to expect, but the model produces the result without assumptions about the macro outcome, only about individual behaviour.

We also see that segregation indices yield the same values across groups either when there are only two groups (irrespective of the share of each in the population, see Figures 4 and 5) or when there are more than two groups and all of them have an equal share (see Figure 9). When proportions are not equal, it is the minority that "controls" how intolerant preferences can become. This becomes clear, for example, in the single and multiple minority cases (Figures 5 and 7), the duopoly (Figure 6) or the ladder (Figure 8). While the indices of the majority group stay low, it is the minority that is forced to "seek each other", as it were, in order to fulfil their preference requirements. This results in their spatial pattern becoming more and more clustered and, at some level of τ, ceasing to meet equilibrium. In this sense, we could call this phenomenon the "rule of minority" and, at first, could appear counter-intuitive, as it is typically assumed that it is majorities that control how segregated an area is. However, once it is put in terms of the behavioural rules underlying the model, it becomes clear that it is what one would expect.

There are three possible cases we can observe in terms of how the indices evolve as τ increases. In some configurations, indices achieve a maximum value at some level of τ and remain at that top thereafter. This sometimes represents the theoretical maximum of the index (e.g. the segregation and isolation measures in the benchmark), although it is not necessarily the case (e.g. Ellison &

Glaeser or Maurel & Sedillot). In other instances the maximum achieved by an index appears to be only limited by the fact that we are forced to restrict the number of trials for computational reasons, and it seems as if the index could potentially continue growing (e.g. the Gini in the single minority). It is possible that, if we were free of computational constraints and thus able to infinitely try initial conditions, we could extrapolate some of the plots, although not others, since it is potentially not always the case that a solution exists for certain conditions. Yet in other contexts, the indices achieve a global maximum and, if preferences continue turning even more restrictive, their values decrease. This is the case of the Gini coefficient but, more notably, of the concentration measures which, almost consistently, achieve a peak at a medium value of τ and then decrease.

Besides these global trends, we can also examine the behaviour of particular indices and their relation with τ. Both segregation measures, the isolation index and the Gini coefficient (except for the duopoly and ladder cases) are monotonic on τ. This is a desirable property because it allows the researcher to interpret higher values of the index as a higher "revealed" preference for co-location with members of the own group. Oppositely, the two measures of concentration (EG_{gt} and MS_{gt}) do not satisfy this characteristic in most cases, and this raises issues about their suitability to obtain a clear view of the distribution and arrangement of the groups (e.g. the same value could be representing two entirely different maps). We conclude that concentration indices should be used and interpreted with much caution when measuring spatial diversity.

Finally, a related but slightly more subtle feature of an index is the form of the relationship with τ. A rather linear function appears more useful because it allows us to compare and track changes more clearly than if the same quantitative increment implies different changes depending on at which level it occurs. In this context, an index that only varies within a very small range of values of τ (i.e. moves from its lowest value to its peak very quickly) is less useful than one for which change happens over a larger range of τ values. Although it is complicated to rank the analysed indices along these lines, it becomes clear from Figures 4-9 that this property very much depends on the scenario considered and the particular group. Both segregation indices perform very similarly, as isolation does too, tending to have a specific range where most of the change happens. The benefit of the Gini coefficient in this context is that it displays similar behaviour but, in most of the configurations considered, significantly varies along a larger range of τ values.

6   Concluding remarks

Cultural diversity has over the past years become a fashionable term, but its actual and reliable measurement is still fraught with many problems. A first major concern emerges from the fact that there is a great diversity in cultural diversity measures, each with its own meaning and interpretation. In various cases, these measures have a normative framing, as their content is based on stylised or value-loaded meanings 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 indications has also been developed. Clearly, the findings from a range of different cultural diversity measures will by necessity 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

segregation and isolation measures of spatial diversity are shown to be non-decreasing in increasing preference for within-group co-location, but the Gini coefficient and concentration measures are not.

In the effort to keep the theoretical framework as simple as possible, we have been forced to make several simplifications. This is particularly true for the version of the Schelling model we have used to generate maps. 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 maintaining the general character of this approach. Some of the most direct additions include more sophisticated behaviour rules that account for several attributes such as: labor markets, house prices, different generations and age or mobility costs. Some of them have already been implemented, as mentioned in the literature review, but accounting for them when comparing measures of diversity could yield new insights. Another complementary addition would go in the direction of blending the MC-ABM approach outlined here with microsimulation methods to explore its forecasting properties in real-world settings, with an eye towards its usefulness to better understand potential policy measures. These are all very promising future extensions of this work.

In general, the conventional Schelling results on cultural segregation are confirmed with our MC-ABM experiments. However, it is noteworthy that the above mentioned caveats are present in a purely spatial context. If we were to introduce time as a dimension in the analysis, an even more pronounced situation would likely arise. Consequently, the mathematical-statistical underpinning of spatial dynamics in cultural diversity remains still a scientific puzzle. Although we have shed some light, there continues to be a clear need for robust summary indicators of the spatio-temporal evolution of cultural diversity.

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: A.J. Heppenstall et al. (eds.) *Agent-Based Models of Geographical Systems*. Springer.

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

Brown, Z.A. and Moore, R.D. (1970). The intra-urban migration process: a perspective. *Geografiska Annaler*, B(52): 1–13.

Chen, K., Irwin, E.G., Jayaprakash, C., and Warren, K. (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, W.A. (1991). Residential preferences and neighborhood racial segregation: A test of the Schelling segregation model. *Demography*, 28(1): 1–19.

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

Clark, W.A. 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, A.T. (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, D.M., Glaeser, E.L. and Vigdor, J.L. (1999). The rise and decline of the American ghetto. *Journal of Political Economy*, 107(3): 455-506.

Duncan, O.D. and Duncan, B. (1955). Residential distribution and occupational stratification. *American Journal of Sociology*, 60(5): 493-503.

Ellison, G. and Glaeser, E.L. (1997). Geographic concentration in U.S. manufacturing industries: a dartboard approach. *Journal of Political Economy*, 105(5): 889-927.

Feitosa, F.F., Le, Q.B., and Vlek, P.L. (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 Waren, W. (2005). Overlooked implications of ethnic preferences for residential segregation in agent-based models. *Urban Studies*, 42(11): 1893–1917.

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

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

Maré, D.C. (2005). Concentration, specialisation and agglomeration of firms in New Zealand. Motu Working Paper 05-12, Motu Economic and Public Policy Research, Wellington.

Maré, D.C., Pinkerton, R.M., Poot, J. and Coleman, A. (2012). Residential sorting across Auckland neighbourhoods. *New Zealand Population Review*, 38: 25-54.

Massey, D.S. and Denton, N.A. (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., Poot, J. and Sahin, M. (eds) (2012). *Migration Impact Assessment: New Horizons*. Cheltenham UK: 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*, 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, S.F. and O'Sullivan, D. (2004). Measures of spatial segregation. *Sociological Methodology*, 34: 121-132.

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

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

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

Schelling, T.C. (1978). *Micromotives and Macrobehaviour*. Norton: New York.

Simon, J.L. (1999). *The Economic Consequences of Immigration*. 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.

Urry, J. (2000). *Sociology Beyond Societies: Mobilities for the Twenty-first Century*. International Library of Sociology. Routledge.

van Mourik, A., Poot, J. and Siegers, J.J. (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.

Table 1: Diversity measures

	Map 2 (a)		Map 2 (b)	
	Group 1	Group 2	Group 1	Group 2
Segregation	0.1021	0.1021	1.0000	1.0000
Modified Segregation	0.0510	0.0510	0.5000	0.5000
Isolation	1.0240	1.0240	1.9995	2.0005
Gini	0.5813	0.5925	0.6999	0.7542
Ellison & Glaeser	-0.0005	-0.0005	0.0016	0.0016
Maurel & Sedillot	-0.0022	0.0013	-0.0080	0.0111
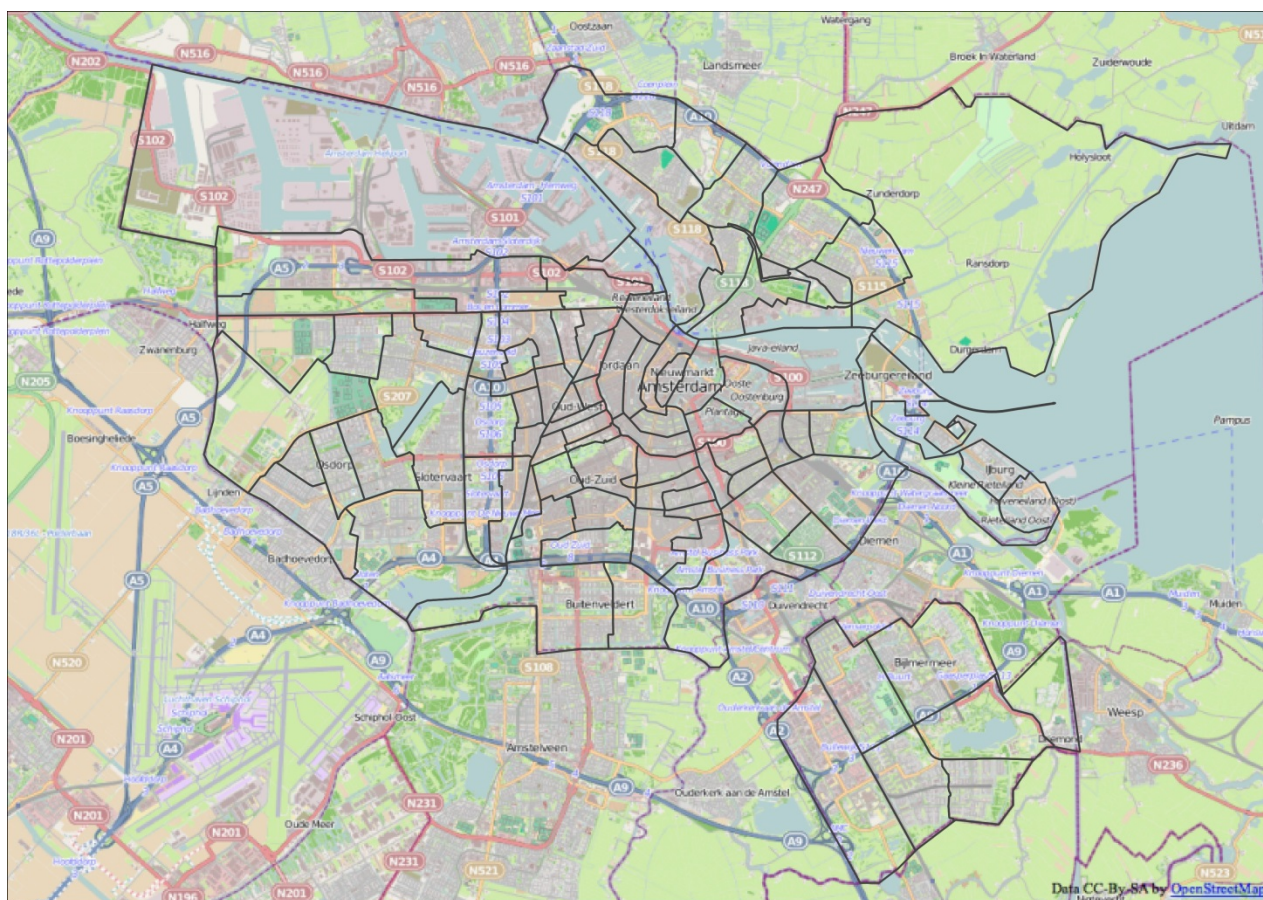
Figure 1: Amsterdam neighbourhoods
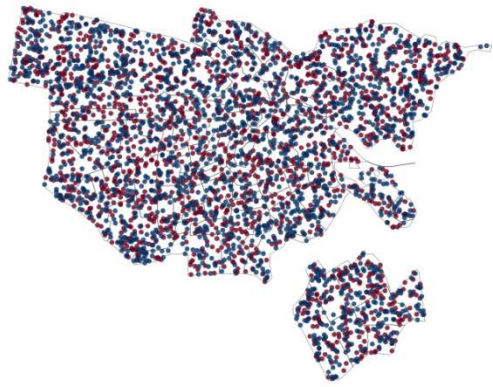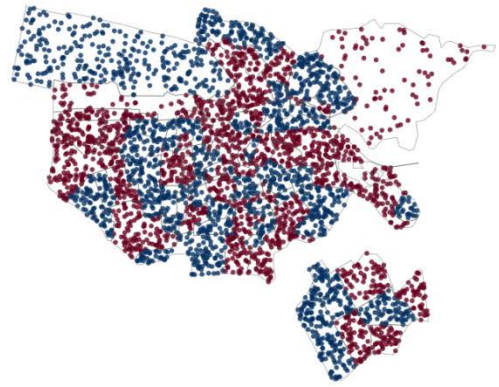
(a) τ = 0.   (b) τ = 0.65

Figure 2: Output map examples (two groups)
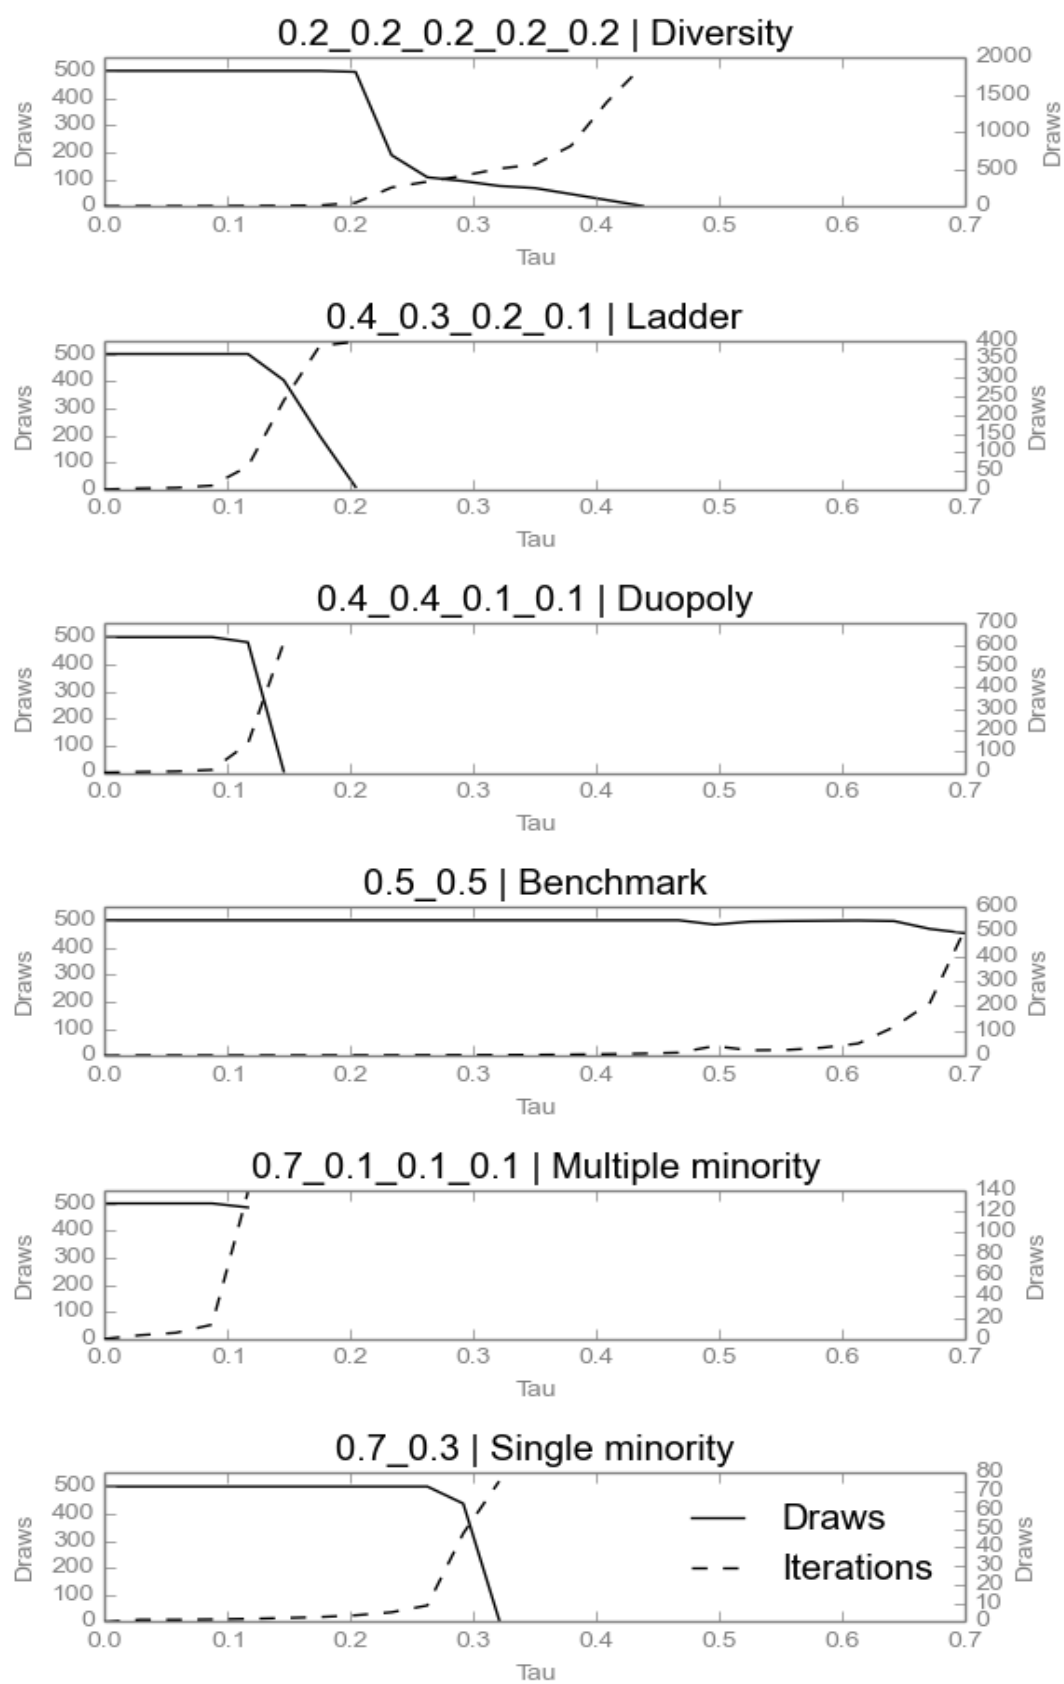
Figure 3: Simulation descriptives
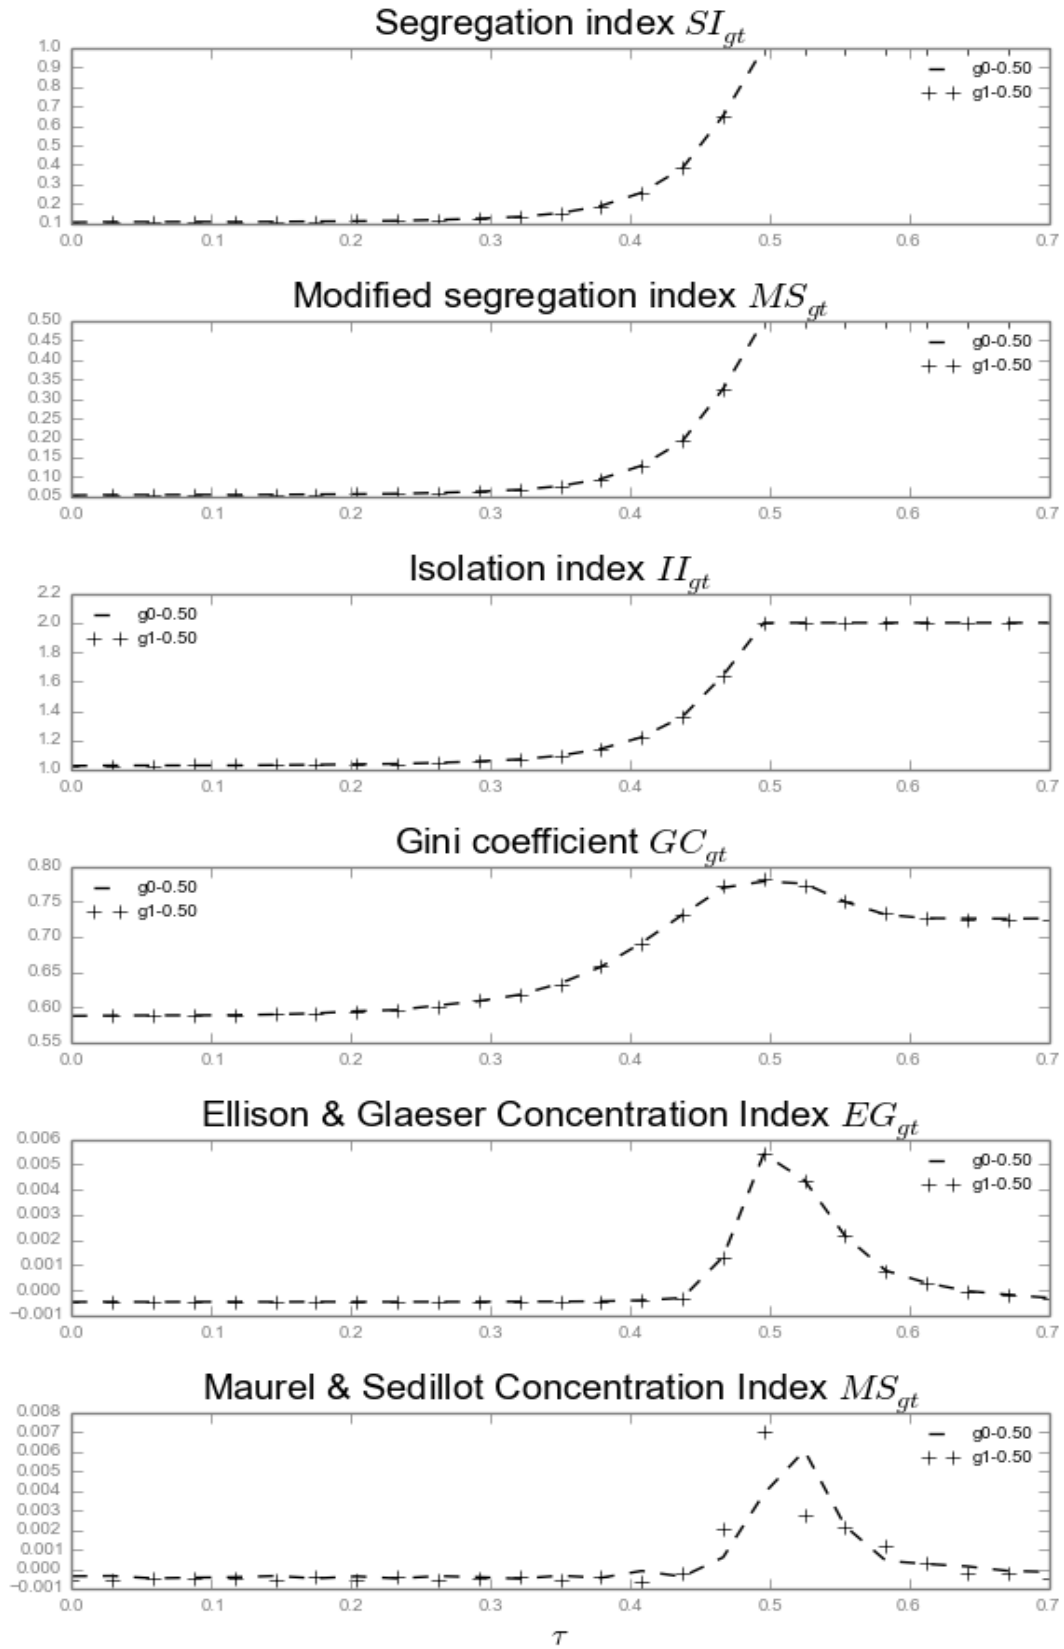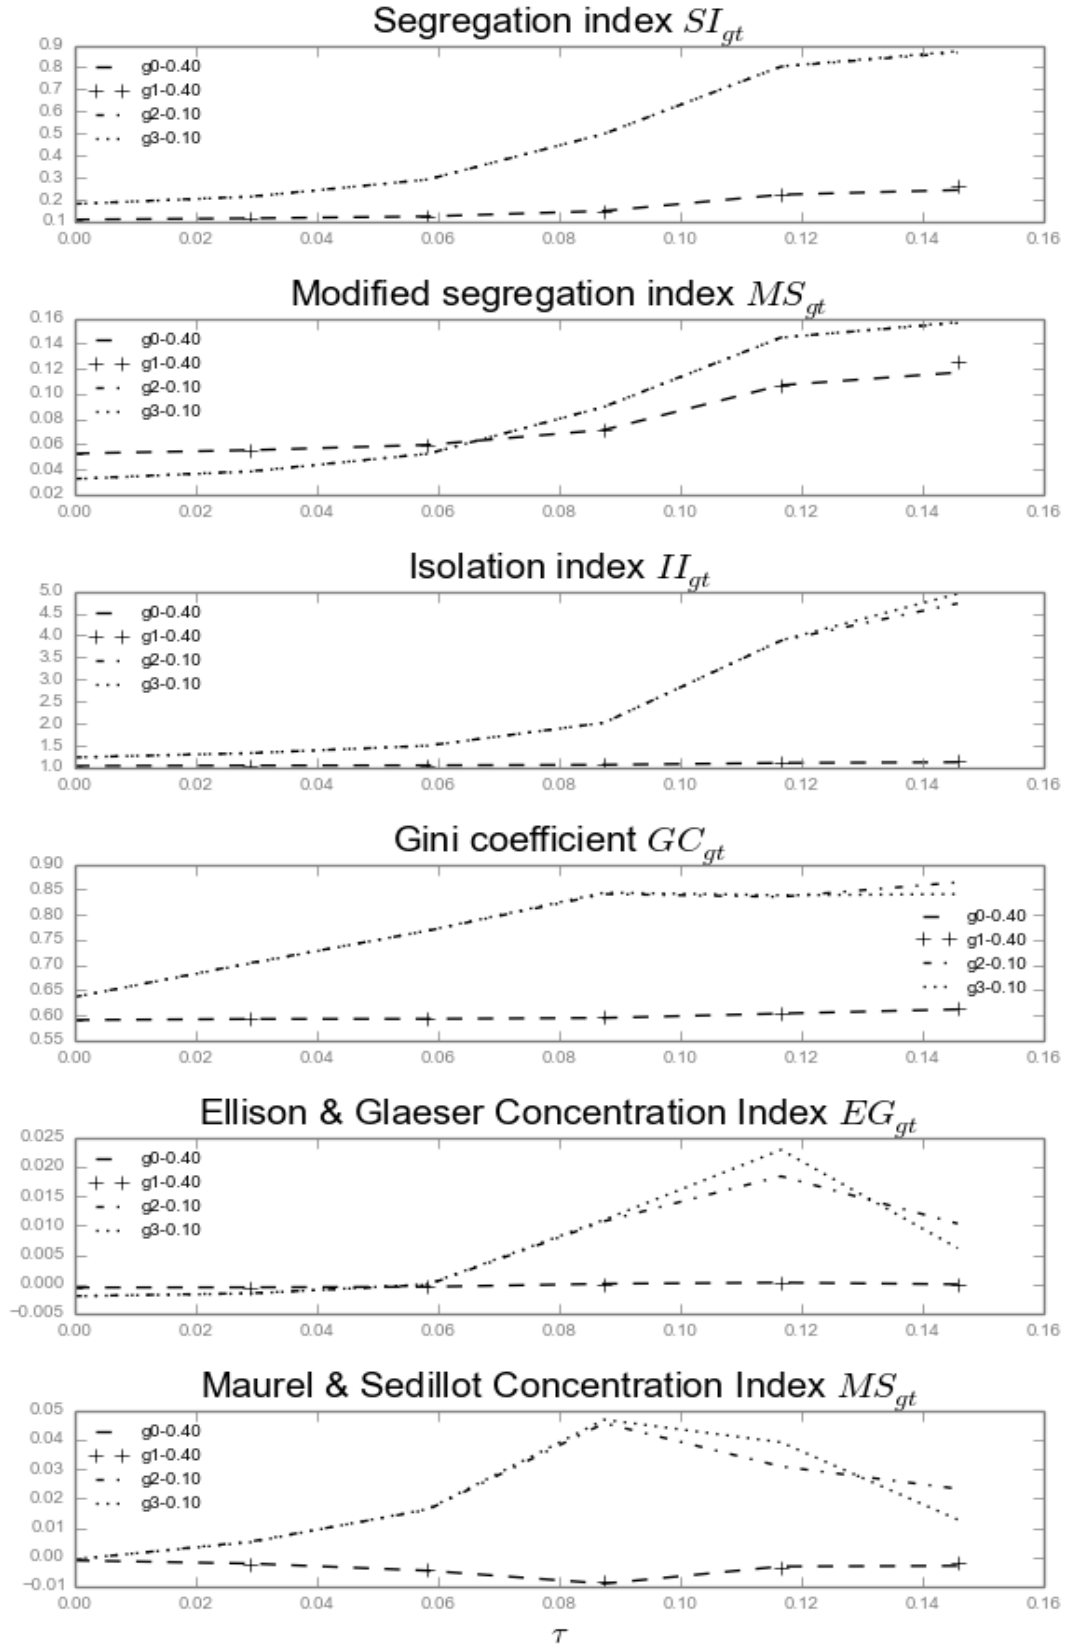
Figure 4: Benchmark
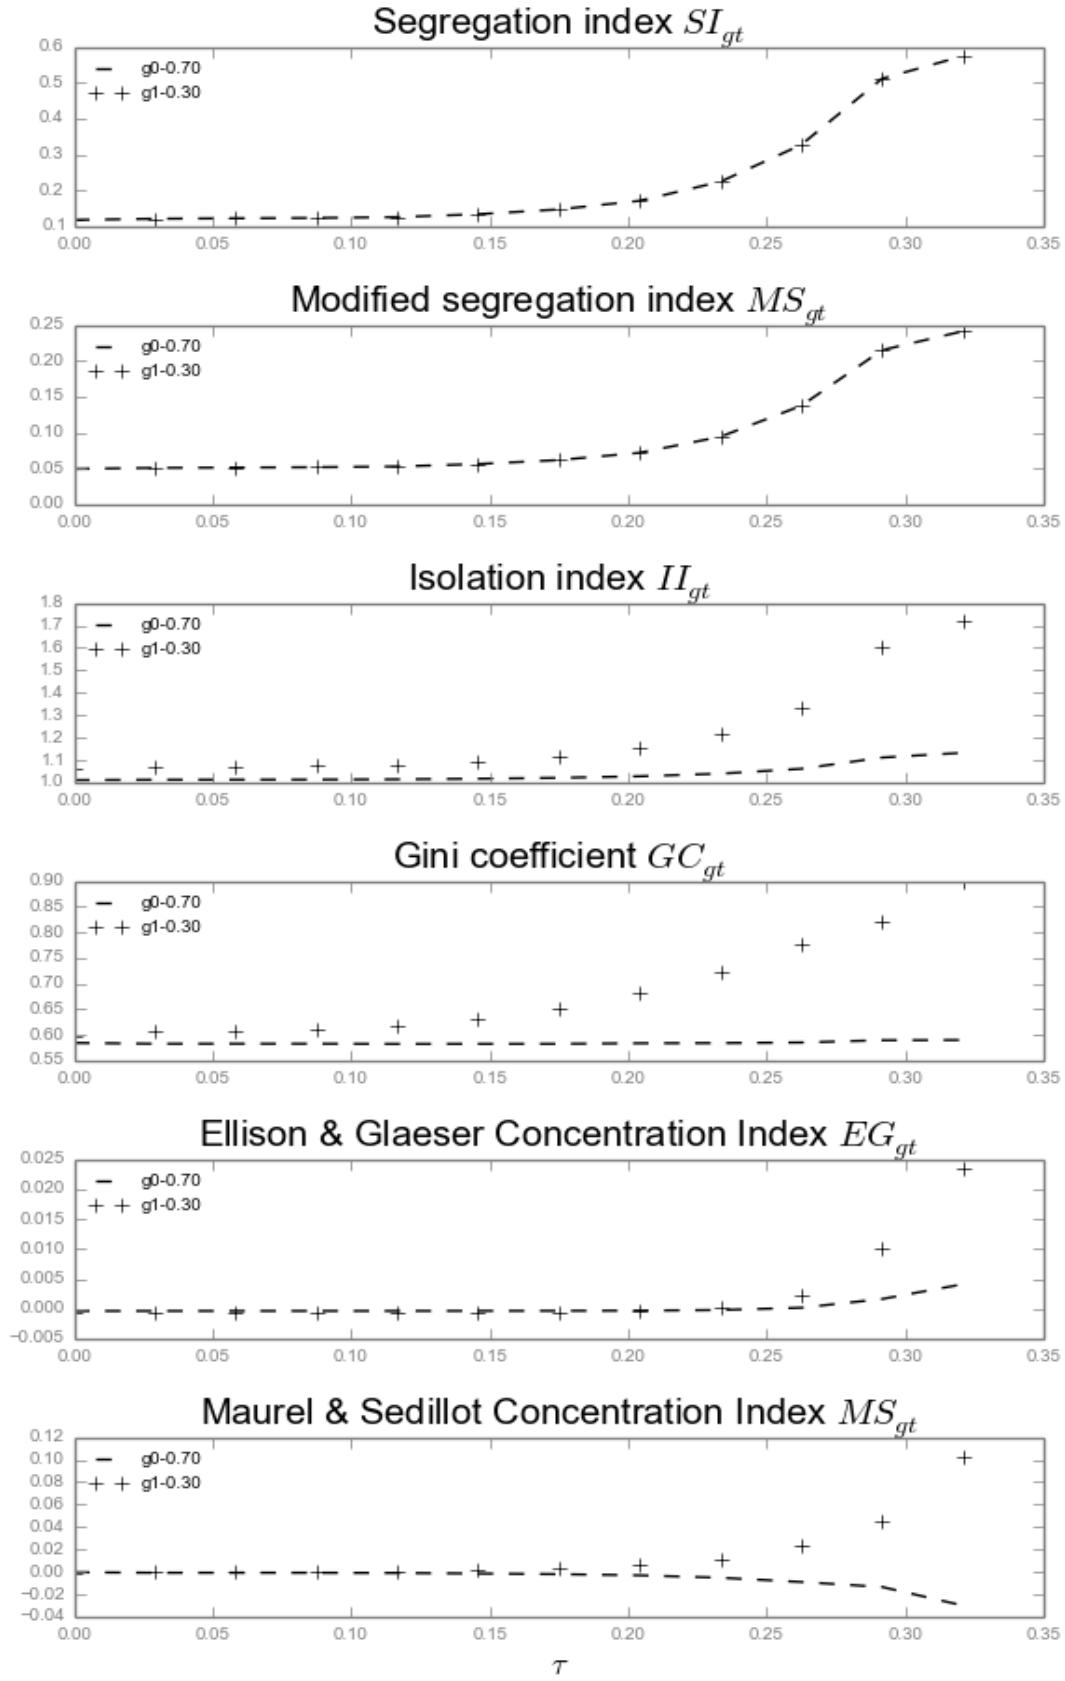
Figure 5: Single minority
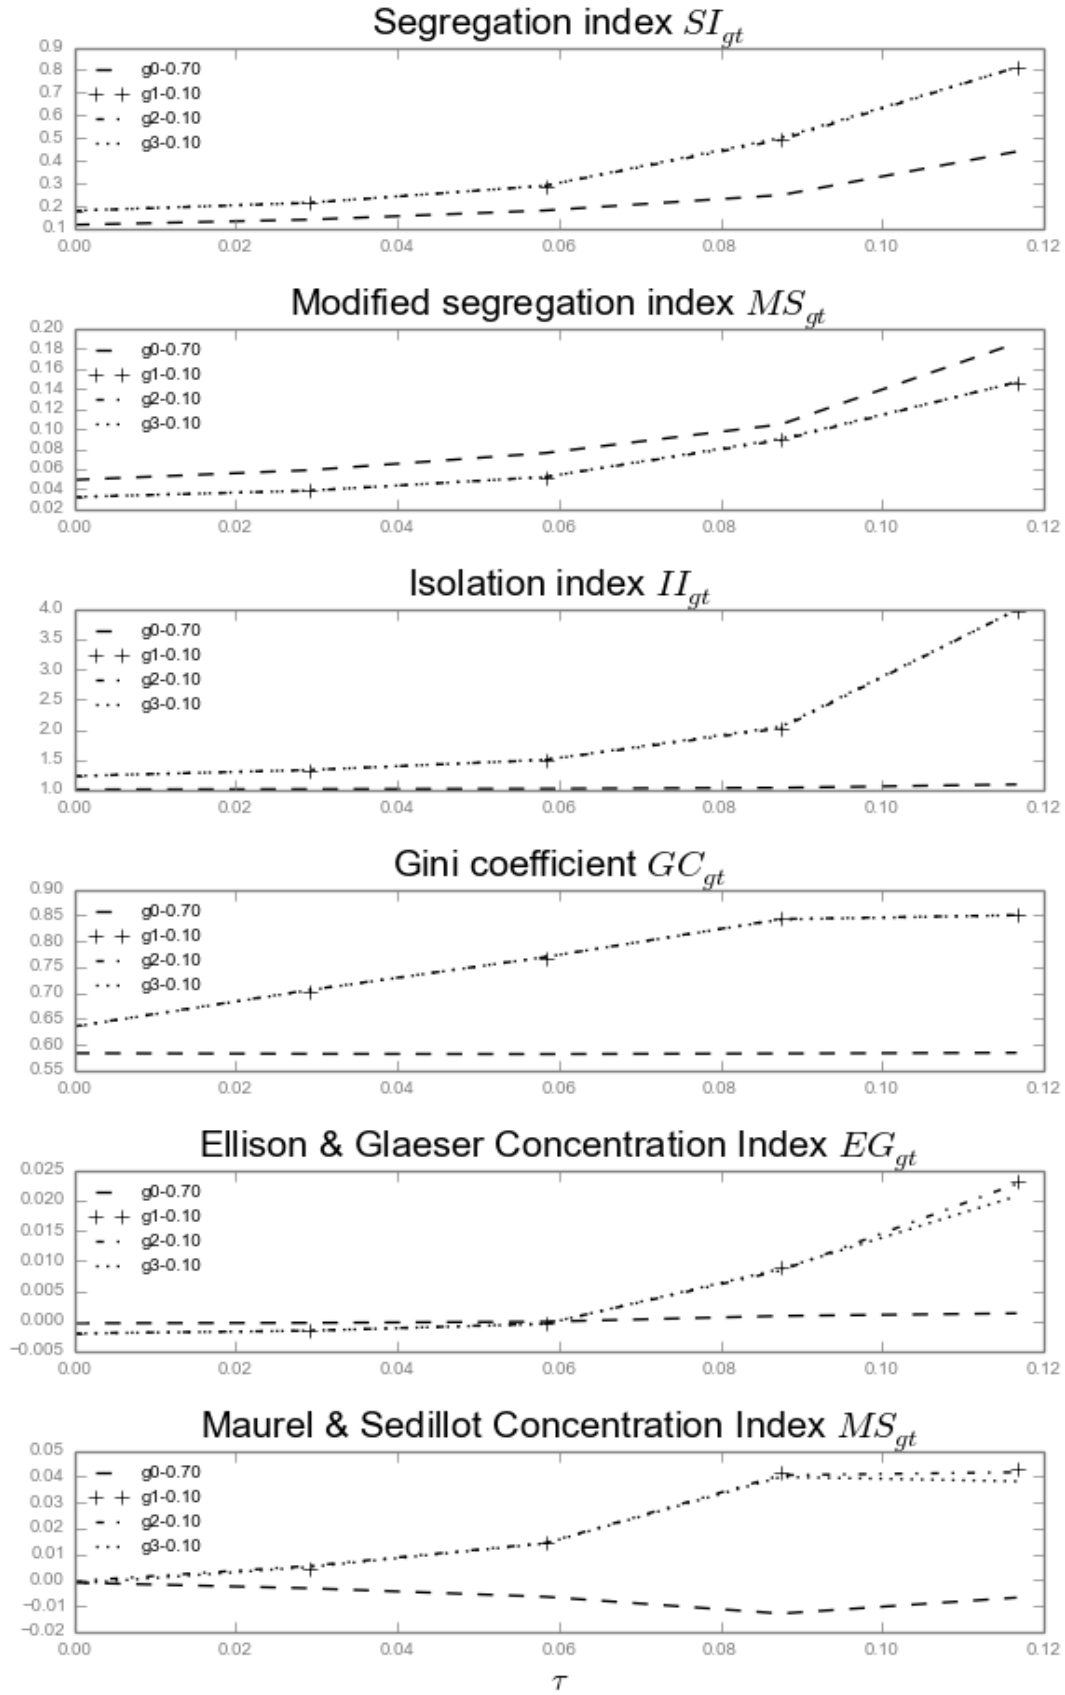
22

Figure 6: Duopoly with 2 minorities

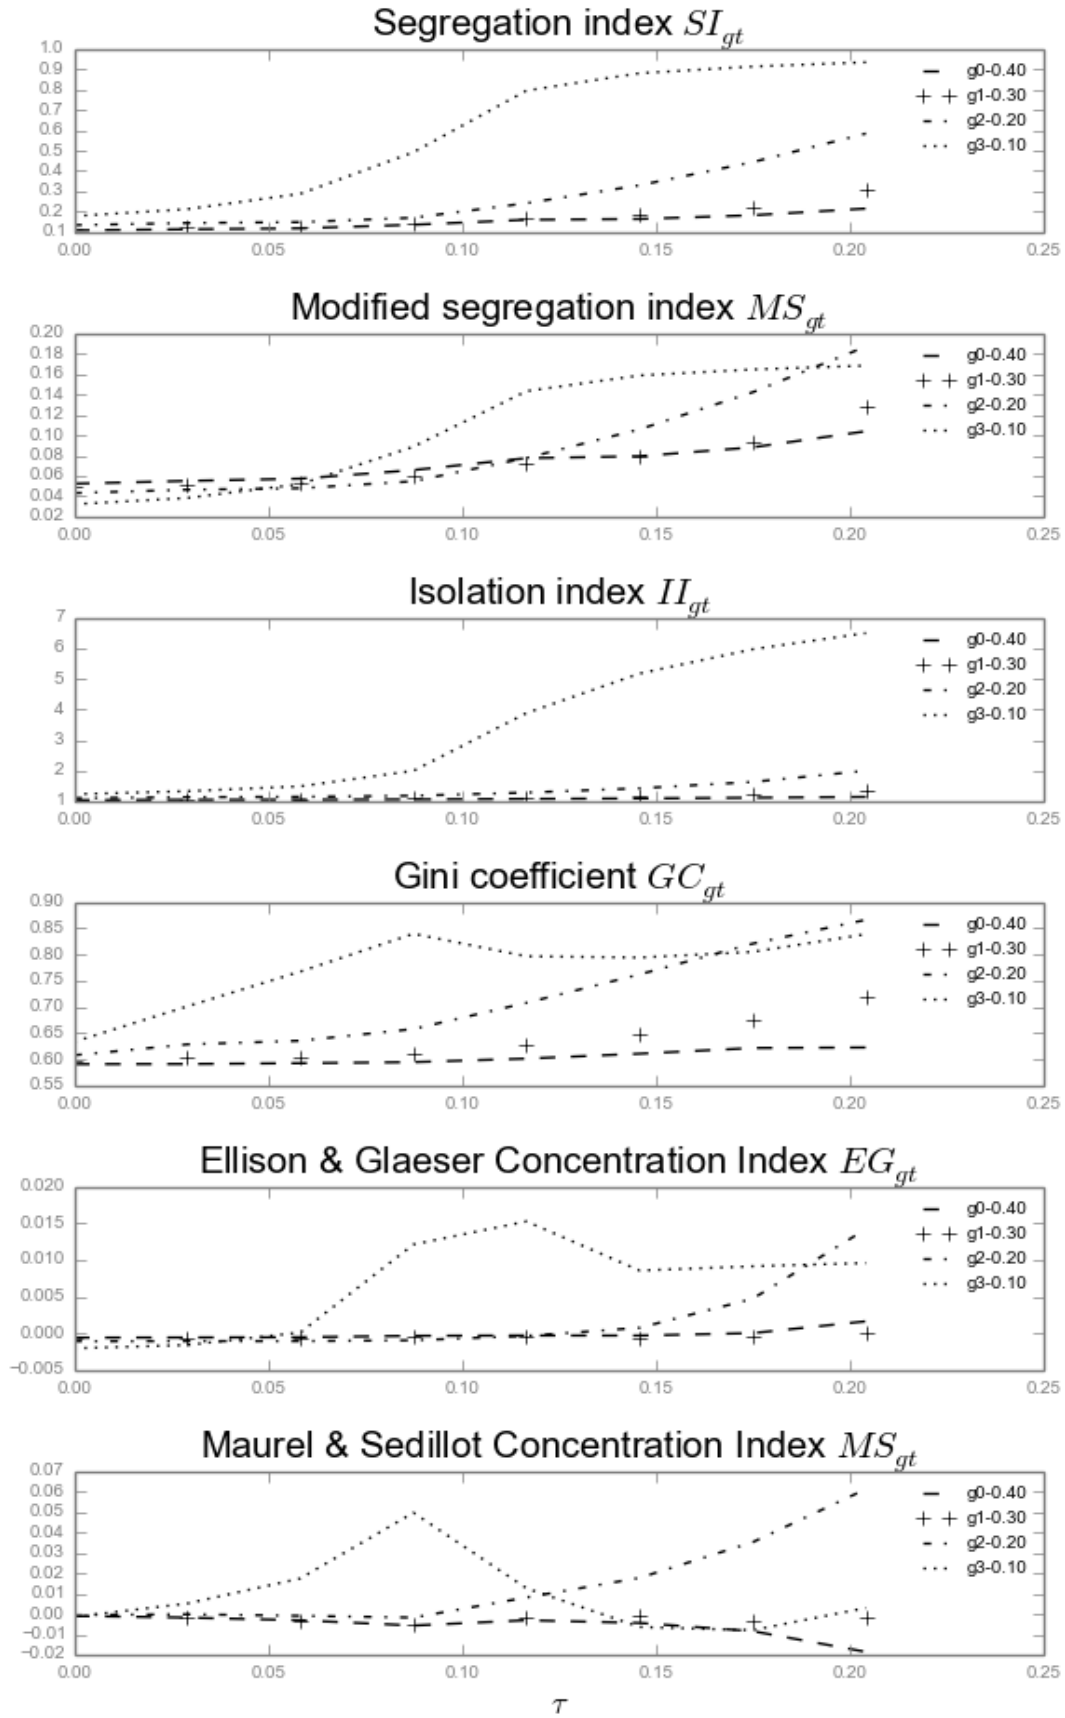23

Figure 7: Multiple minority

24

Figure 8: Ladder

25

Figure 9: Diversity

26