Agent-based modelling of complex systems in political science: Social norms and tolerance in immigrant societies

Linda Urselmans

A thesis submitted for the degree of Doctor of Philosophy

Department of Government

University of Essex

March 11, 2018

Contents

List of Figures

List of Tables

To my brother

Abstract

Migration is a political issue that has received more attention in recent years. Many questions remain as to how Western societies can successfully absorb migrants-economic arguments have largely been in favour of migration, but the social impact of diversity in previously homogeneous societies has been subject to ongoing debates in social science.

Migrant societies are complex social systems with many interacting moving parts. How do rapid migration-changes in society affect the hosts? How do norms of tolerance towards minorities hold up when intergroup conflicts emerge? Can segregating behaviour of different population groups be reduced by encouraging different settlement locations for new migrants? The questions address both the physical aspect of migrants entering an already populated space, and the social dimension in which the hosts are adapting their attitudes.

I develop a Schelling model using Agent-based modelling to address these questions. I introduce the concept of external migration into an existing society and test how, by varying the kind of migration, introducing diversity affects local tolerance. In the second chapter, I show that large-scale migration results in short-term shocks to the populace, but that these effects are heavily dependent on the population density and how large the native majority is. In Chapter 3 I implement a version of the 'contact hypothesis' which stipulates that contact with out-group members increases tolerance and I show that the adaptability increases the importance of native majorities further. In the fourth chapter, I move on to the social norms of tolerance, introducing an ABM in which agents can deceive others by signalling false information about their true attitudes. I show that the emergent pattern of these behaviours can lead to a false consensus effect in which the perceived majority public opinion is unstable.

The thesis is able to generate societies that bear many similarities with the Western countries of today and can suggest explanations for the mechanisms that lead to changes in public opinion more negative towards migration, as well as reasons for growing separation of different population groups.

Acknowledgements

I am indebted to my two supervisors Hugh Ward and Steve Phelps who have guided me on this journey. Hugh, your teaching gave me the impulse to pursue a PhD in the first place. Without your intellectual and emotional support, this thesis would not have been possible. I will miss our discussions about environmentalism in video games, history of tandems, and every topic in between. I admire your intellectual curiosity and breadth of knowledge.

Steve, thank you so much for your continued support and for teaching me to not be afraid of anything that can be studied and understood. You gave me confidence to believe in my skills, to expand my comfort zone, and to seek answers to questions that I had asked myself for a long time.

To my mum, who had to endure all these long Skype sessions about migration, agents, and programming. My sometimes futile attempts to translate my work to you were so important. Your interest in my interests sustains me. To my dad: your words may be few, but they carry great weight. When everything is in turmoil, you are the solid rock. To my brother: I know you like a grumpy attitude, but your desire to support me and help me with your expertise really touched me. I am so proud to have you as a brother. To my sister: for all these years that I pursued my dreams, I could always come back to you and feel at home. You just have that little something, the joy to share and to help- as a little sister, this is something I will always look up to.

I am also grateful to the departmental admin team that has provided me with so much support and assistance. A special mention to you, Nicola!

Thank you Aziz for all the hours of discussions that we shared. They have really expanded my horizons on many issues. Rachel, you were my backbone for all these PhD years and I am so grateful for having had the pleasure to meet you.

I thank the Economic and Social Research Council UK for providing me with the funds that enabled my PhD in the first place.

1 Introduction

Computer models as a method for scientists have enjoyed widespread use in the natural sciences since the advent of modern computing. Agent-based models (ABMs) are computer models that allow the construction of autonomous individuals and groups, and their complex interactions. These models can complement the commonly used methods in Political Science. They can be deployed to test existing social theories and mechanisms where empirical research has been inconclusive or where social theories can't be approached with analytical models.

In recent decades, social scientists too have begun to adapt computer models to test social theories and to simulate *artificial societies* (Epstein, 2006). Agent-based models are a class of computer models that focus on simulating autonomous individuals, groups and their interactions within a system. ABMs can be deployed as an alternative method to equation-based modelling. While there has been an increase in computer modelling in fields such as behavioural economics, they remain underused in the social sciences (Helbing, 2012). Bruch and Atwell (2015) lament its "minimal impact on mainstream sociological research", suggesting that a lack of dialogue between ABM work and data-driven work is one of the culprits.

One of the areas of research which has both an active ABM and a data-driven community, is that of migration. Agent-based modelling approaches can be useful to study migration, which spans several disciplines. Migration is a complex system of decisions, opportunities and costs. The underlying scientific approach of agent-based modelling is rooted in complexity science, which will be briefly discussed below.

Complexity theories aim at understanding the properties of complex systems. There is no single complexity theory, in large part because there is no single field of complexity science. Instead, disciplines ranging from physics, biology, economics, anthropology and sociology have applied versions of complexity theory to systems of interest. In Political Science, the complex systems of interest can be voters, parties, elites, social groups or entire societies.

Complexity and its intricacies have been studied for several centuries. In 1814, Pierre Simon Laplace posited that, if one had all the information of any given system (say, the whole universe), it would be possible to predict everything for eternity (Mitchell, 2009). This notion was coined the 'clockwork-universe': a universe made up of lots of parts, wound up and working according to Newtonian laws (Bryne, 1998). Laplace and others thought that the knowledge of every particle and velocity in the universe was a practical restriction: doable in principle but impossible in practice (Mitchell, 2009). With this mindset, researchers likened societies of humans to clocks as well: a complex (mechanical) system that is made of many individual parts which, in interaction, form what we know as a clock (Sawyer, 2005). According to this view, if we knew every human being and every interaction in society, we could predict its future. This view was challenged in subsequent developments in physics, but the approach remains useful. Whilst we may not understand, explain and predict everything, we can try and explain part of the systems of interest.

There are many complex systems of which we understand parts, but not all: Insect colonies, the brain, the immune system, economies and the internet (Mitchell, 2009). Economies consist of people and companies that engage in buying and selling behaviour. Some complex systems respond to the external environment and adapt to it through means such as learning. For example, economies reside in different political and natural environments. The economy as a system comprises trading elements, but the environment can affect the behaviour of individuals and groups within economies and vice versa. Political changes can result in changes in regulation, thus directly altering the nature of transactions in parts of the economy. Climate determines what type of fruit can grow on which part of the planet, thereby influencing the kind of economies that exist in each of these regions. Conversely, excessive gathering of sources such as large-scale deforestation or deep mining can impact the ecosystem. Individuals that form part of an economy can also exert influence on the political system that coexists, and potentially achieve changes in politics that in turn influence the economy (Mitchell, 2009). All these moving parts and interactions influence the complex system of 'the economy'.

To try and understand the nature of complexity helps our understanding of its moving parts. If society does not function as a clockwork, how does it work? Once society is interpreted as such, it raises the question of how these societies can come into existence in the first place. If there is no super-structure or a concious will to build a society, how can it exist in the first place? What makes some societies withstand internal and external pressures, when others crumble?

These answers invite investigation and cannot be answered by studying idealized problems of macro-states. This is the point in time in which the advent in computing power opens up opportunities previously barred to scientists. The ability to carry out millions of computations at a fraction of human time allows for the construction of complexity in the virtual laboratory. To construct society, all of its inhabitants, the agents, are created. The system is the collective of these agents. A computer model that constructs a system from the 'bottom-up' in which the agents determine its output, is called an agent-based model.

Agent based modelling adopts many of the principles of complexity theory. To understand the whole (the macro-level), agent-based models construct the parts (the micro-level) that constitute the whole and the interactions between the parts.

The aim of this thesis is two-fold. Firstly, it addresses two research questions to understand the impact of migration on societies and the importance of social norms that shape expectations and attitudes. Secondly, it demonstrates the various applications of agent-based modelling that can be applied in Political Science. Each chapter will address different elements of the subject and method matter. Chapter 2 focuses on the physical aspects of migration: how does the size and frequency of migration differ in its impact on the hosts, and what effect do different settlement locations have? How does the size of the native and migrant group interact with these effects? The chapter will demonstrate how building on an established ABM can make additions easier to understand. I use stylistic facts and existing models to guide the model design. Chapter 3 builds on the physical elements and incorporates social theory in the form of adaptive tolerance. Given the native share of the population, how do natives and migrants adapt their tolerance in reaction to migration? The ABM in this chapter directly implements an existing social theory to test its premises. Chapter 4 focuses on the social norms that influence perception of migrant acceptance and the need to appear tolerant. The chapter demonstrates the use of agent-based modelling to contribute to theory development by testing existing elements of the theory.

The remainder of this introduction is structured as follows: Section 1.1 provides a literature review of Agent-based modelling and relevant work in the social sciences. I introduce *boids*, a well-known ABM, to demonstrate the strength of the method; followed by a discussion of the weaknesses and pitfalls of agent-based modelling. The following section reviews the relevant subject literature on migration, diversity, tolerance and public opinion. Lastly, the research questions are outlined.

1.1 Agent-based Modelling

Socio-economic systems have been notoriously hard to model (Helbing, 2010). They are more complex than physical models, providing the challenge that simple models run danger of being too abstract, yet more realistic models may be too complicated to be fully understood, thereby defeating the purpose of modelling: providing a simpler version of the real version in order to understand or predict parts of, or an entire system. Many different types of models exist that tackle socio-economic complexity from different perspectives. Each method has different merits. Classical models for instance often implement analytical approaches and can provide predictions of future behaviour under the conditions studied (Gilbert, 2007).

Agent-based models belong to the larger group of computer models, specifically Multi-Agent Systems (MAS) (Mitchell, 2009). These kind of models approach the target system from a Systems theory perspective: to understand the system, the idea is to model the components from the bottom up (Gilbert and Conte, 1995). The study of real-word societies (such as the rise and fall of the Mayan culture) was among the first objectives of social science computer simulations (Gilbert and Conte, 1995). Agent-based modelling is used to provide the micro-foundations of changes observed on macroscopic levels (Epstein, 2011). ABMs generally seek to provide explanations of underlying processes as opposed to describing observed macro-level outcomes (Smith and Conrey, 2007). In the social sciences, the equation-based model has gained great popularity (Gilbert and Conte, 1995). These mathematical models are usually employed in empirical work; structural equation models for example test the relationship between explanatory variables, usually part of a social theory. These models are measured by the goodness of fit with the supplied data, and the focus lies on the description of the relationship between the key variables. A statistical approach to segregation might be for example finding that whenever ethnic minorities constitute a certain percentage of the overall population, segregation rises. Finding such a regularity (if it exists) is useful to describe and predict the phenomena, but doesn't look at why the phenomenon emerges in the first place. Agent-based modelling emphasises this aspect of modelling: generating possible explanations of underlying processes and showing how a phenomenon might emerge. Instead of focusing on 'what happened?', the focus lies on the set of preconditions that exist before something happened. This is very useful for testing and advancing existing social theory. Theorists have had problems specifying social theories to a sufficient degree, and a wide gap between some areas of theory and empirics exists (Gilbert and Conte, 1995).

Because of its experimental nature and lack of empirical restraints, computer simulation can test sets of conditions surrounding a theory, thereby informing future debate in theory development. The many microeconomic models developed in economics adopt a 'representative agent' framework (Epstein, 2011) which assumes that the collectivity of agents is represented by the aggregation of individual decisions. The population of agents is assumed to be homogeneous. Agent-based models usually feature sets of heterogeneous agents, instead focussing on the dynamic interactions between them. Agents' behaviours can be internally driven through intrinsic attributes, or externally driven through the interaction with the environment and/or other agents (Epstein, 2011).

Computer modelling has brought the advent of 'generative' science: the ability to artificially re-create and run experiments gives researchers the ability to generate a set of explanations and proposed relationships (Epstein, 2006). Experiments are designed to isolate the target system from its environment, and to subject the system to a number of conditions in this controlled environment, so that the result can be attributed to the experimental parameters. In social science, experiments usually involve people- individuals or groups- and even in isolated environments, there are a number of cognitive and physiological elements that cannot be controlled for or have unintentionally not been controlled for. This makes the testing of some social theories difficult. Especially when individual-level and group-level behaviour are the subject of interest, it is often difficult to capture this link. One way of overcoming the micro-macro gap is to collect data from both levels and combine them in multi-level models. The drawback of this approach is that the model is reliant on the quality and quantity of the data. Computer simulations circumvent the experimental problems of isolating people from their cognitive biases because every agent rule must be explicitly modelled. Cognitive biases of agents are the result of a concious decision by the modeller. However, code does not always behave the way that the modeller intends it to behave. Biases can be introduced through simple coding errors. Because computer models are not subject to the same time and money constraints of experimental models, they can test a great range of experimental conditions, including those that would be infeasible or unethical to test in human experiments. Nigel Gilbert has termed this approach ‘exploratory simulation’ (Gilbert, 2007, p.4). A computer simulation can be thought of as an experiment, written as a software rather than conducted in a laboratory.

Agent-based modelling is a type of computer simulation that can use the principles of Object-Oriented Programming(Gilbert, 2007): software that consists of objects, which can have attributes and execute functions. In agent-based modelling, the objects are agents that represent a social entity. In most cases, an agent is an individual; but agents can also represent groups, businesses, governments, countries, or more abstract patterns such as language or trade. The objects of an ABM operate in an environment (analogous to the laboratory setting of a real-world experiment) that is also specified in the program. The environment can be spatial, such as a room, a street, a city or a country. It can be a physical representation of space, including obstacles such as stones, walls or deep water. It can also be social space, such as friends, family or the workplace. In recent years, Geographic Information System (GIS) researchers have turned their attention to agent-based models and build models firmly grounded in GIS data of physical environments (Crooks, 2010). ABM environments must not necessarily be spatial. They can also be abstract, such as a social network which is not tied to a spatial setting. In such a network-based model, the distance between agents may be the strength of their friendship rather than a physical distance, and interactions between agents take place regardless of proximity.

1.1.1 An Agent-based model: *boids*

An often cited example of an agent based model is Craig Reynold's *boids*, 'bird-androids' model (Reynolds, 1987). Reynolds attempted to simulate the behaviour of flocks of birds. After examining the physical features and behaviours of bird flocks, a simplified version was devised. The *boids* follow three rules as they move around a 2-dimensional space:

1. avoid collisions with other boids;

2. try to fly at the same velocity of nearby boids

3. attempt to stay close to nearby boids

Each boid has these same rules. As a collective, they behave like real bird flocks (later versions of the model included object avoidance and 3-dimensional space). Figure 1.1 shows screenshots of a basic version of boids called the 'flocking' model

Figure 1.1: Screenshots of flocking behaviour emerging

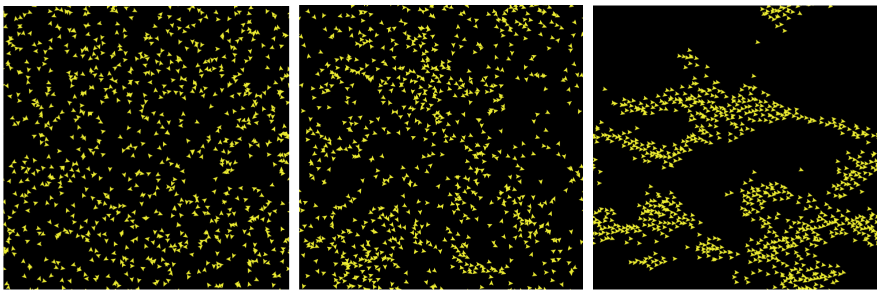

and is part of the NetLogo standard library (Wilensky, 1998). Boids, or birds, are represented by yellow arrows on black background, representing empty space. At the start (left-hand side), all the birds are randomly positioned. Each bird has a unique and randomised position, velocity and direction. As the model progresses (centre screenshot), and the birds act based on their three rules, they start clustering into visible groups (flocking) and start flying in similar directions. After a while (right-hand side), clear bird flocks have formed and they all fly in the same direction.

The *micro level* of this model is the boid. Each individual boid has certain attributes: position and velocity on the grid, and a *set of rules* that each boid follows. A rule in an ABM is designed to guide the behaviour of individual agents (boids). Boids have three rules in this case, but they can be expanded. Boids are also not necessarily homogeneous. They share the same ruleset, but other versions of the model can include *heterogeneous* boids: some follow one ruleset, some follow another. Rules can also be contingent upon certain special conditions: for instance, if a boid finds itself all alone, it will abandon the three-rule approach and will prioritise re-uniting with other boids first, before resuming the three-rule-based behaviour. Boids can also differ in size or shape, impacting their behaviour in the world. Agent rules and behaviour can be deterministic or probabilistic. For example, a rule may stipulate that a boid will always adjust their velocity according to their nearest neighbour. Such a rule could also include a percentage chance of a random change in velocity. The *macro-level* of the flocking model is the flock of boids, or the population of boids. The strength of this model is that only three micro-level rules are required to generate macro-level outcomes (flocking) that mimic real-world bird flocking. This outcome provides a possible explanation for how bird flocks can form as a result of the collective behaviour of individuals. Crucially, flocking is an *emergent* property of the model. Boids are not programmed to form flocks, but rather, to deal only with their immediate surroundings. The collective of these decisions leads to patterns of flocking.

1.1.2 An Agent-based model in Political Science: *Strategic voting*

Agent-based models are not widespread in Political Science, but they do exist. Clough (2007) introduced an Agent-Based model of strategic voting, basing the model on the mechanisms of Duverger's law. Duvergers law posits that the electoral system influences which party systems develop in a democracy. Proportional representation, according to Duverger, is more likely to give rise to multi-party systems, as more parties stand a chance of governing through coalitions. The 'law' generally holds but there exist exceptions and not everyone shares the belief that the underlying mechanisms are true. One criticism of the proposed mechanism has been the availability of public information about the parties. Voters can only penalise parties for not being able to win, if they can make a projection of what the outcome of the future election will be. If voters do not know about the prospects of their preferred candidate winning, they can't discriminate against third parties on the basis that they don't want their vote to go to 'waste'. The availability of information has been recognised as a key part of the mechanisms of Duverger's law, but existing Game Theory-based models have not included the variance of information across the population (Clough, 2007).

In this ABM, agents are heterogeneous and can have different levels of information or access to information. In this model, the individual agents are voters, connected through a social network (Clough, 2007, p. 318). Each round, voters vote and talk to other voters to find out who voter for whom. Thus, their information about the likelihood of a candidate winning is updated. In the next round, voters cast a new vote, adjusting for the new information that they have received. Voters will try to maximise their utility: to cast a meaningful vote that gets their preferred candidate elected. Voters have a preference order over the candidates (Clough, 2007, p. 319). This ABM is more complex compared to the boids model: it includes social networks, a number of parties, expectations that individual voters (agents) have, and an electoral system that determines how votes are allocated and how a candidate can win. The fundamentals remain the same, however: instead of boids, agents are voters. Instead of flying over the map, voters don't move, but rather interact with peers in their network. Just like boids, voters are heterogeneous: together they form a collective group, but each has individual attributes. Boids have different starting positions, velocity and direction, and voters have different candidate preferences, social networks and information access. The heterogeneity of information access and its abundance is the key contribution of this strategic voting ABM. The results find that strategic voting on its own does not always lead to two-party systems in a single-member plurality electoral system (SMP). When poorly informed, voters cannot coordinate effectively on two parties. The size of the neighbourhood that agents 'talk' to to gather information determines their ability to coordinate. In an age of an abundance of opinion polling this might seem less of an issue, but recent election upheavals such as the US presidential election 2016 have highlighted a long-discussed theme in academia: the ideological polarization of information (Spohr, 2017). Rather than lack of access due to not owning a TV, phone or being able to afford a newspaper subscription, lack of information occurs through various 'bubbles' on social media, which can distort perceptions.

This ABM has used the existing theory (Duverger's law) and empirical findings surrounding it as a basis for how the model is constructed, and was able to generate an additional explanation of the underlying mechanisms that may drive the link between the electoral systems and party systems.

1.1.3 Problems with Agent-based modelling

As with any other method, agent-based models have strengths but also weaknesses and pitfalls, which will be reviewed in turn below.

Agent-based models are subjected to tests in order to verify its purpose and to validate its design and results. Verification of agent-based models involves the testing and fixing of the software itself and ensuring that its is error- and bug-free. Model validation tests whether the model match the target system (Gilbert, 2007). If it does, it must be subjected to sensitivity tests to check whether the outcomes, even though they may match the target system in the real world, are not the result of the specific set of parameters. The tractability of ABM outcomes has been subject to some discussion in the scholarly literature. There are at least two commonly-discussed concerns relating to tractability of agent-based models: i) complexity and uncertainty and ii) technical limitations. Debates over the theoretical assumptions (iii) of agent-based models have been less common in comparison. I discuss each of the concerns of tractability and theory in turn below.

(i) Uncertainty

A major concern for some scholars has been the uncertainty of what drives the results of an agent-based model. Generating outcomes is relatively easy compared to analysing these outcomes, and more specifically, the mechanisms that have given rise to these outcomes. In the past, the lack of computing power has meant that many models were simplistic in their design in order to reduce the resources needed to run these models (Sawyer, 2005). For most social models, this is no longer a concern. This has encouraged researchers to add more features to the model and increase the number of parameters in a bid to achieve greater social realism (O'Sullivan and Haklay, 2000). Epstein and Axtell (1996) published their 'Sugarscape' model which introduced the concept of artificial societies to test social phenomena to a wider audience (Sawyer, 2005). In its simplest version, Sugarscape is a 2-dimensional lattice on which there exist two mountains of sugar. Agents populate this lattice and consume the sugar, which grows back at a certain rate. In *Growing Artificial Societies* (Epstein and Axtell, 1996), the model becomes increasingly complex: agents metabolise sugar at different rates; fertility rates, sexual reproduction, pollution levels and death rates are introduced; agents engage in combat with one another; social networks are introduced and different types of social networks are tested. Sugarscape has been an important cornerstone in the development of Multi-Agent Systems for social simulation (Sawyer, 2005), but some scholars have questioned the effectiveness of countless additions to an existing model (O'Sullivan and Haklay, 2000). With more moving parts, it becomes harder to attribute the model outcomes to causal mechanisms within the model. The model of a complex system becomes so complex that it no longer serves its primary function: to provide a simplified version of the complex reality so that it can be better understood. The growing complexity of models is in part due to the growing capacity of computers to handle complexity in models. If tractability of outcomes can be established, the complexity of a model itself does not warrant dismissal (Sawyer, 2005). ABMs do not abandon the principle of parsimony- far from it. The goal is to generate *sufficient* explanations for observed phenomena (Epstein, 2006). Whilst many more elements in a model may be necessary to explain a system, the set of explanations that is sufficient is desired over those with excess moving parts.

An important difference has to be emphasised: the complexity of the target system, approached by complexity sciences, and the complexity of the model that seeks to study the target system. The temptation to add more variables to a model is not a problem restricted to agent-based models. A regression equation can be extended with countless explanatory variables in an attempt to increase the explanatory power of the statistical model. Model completeness and model parsimony can tackle different kinds of questions. A parsimonious model may be useful in solving for a particular problem in idealized circumstances- more complex, more complete empirical models may be used to replicate and/or predict an existing problem such that model outcomes can be turned into policy. The baseline prisoners dilemma gives us valuable insight into fundamental problems of cooperation, but offers little insight into specific problems. A more complex model of the prisoners dilemma that accounts for more explanatory variables can thus provide a different angle. The same is true of agent-based models. Their strength lies in the ability to model the micro-states of complex systems.

Complex systems are non-linear and cannot be reduced to their micro-states. To criticise agent-based models for their complexity misses the different focus that the method offers. Agent-based models are subjected to scientific scrutiny and their tractability should be ensured.

(ii) Technical limitations

The model input of an ABM is determined by the modeller. Which parameters, what kind and their respective range should be informed by existing theoretical and/or empirical findings; as should the agent rules and the environment that agents reside in. When model parameters are informed by stylised facts, there is a danger to assume that only this set of parameters can serve as a potential explanation for the subject of study. That is, if one explanation is found, it may be assumed that this is the only explanation that exists for a particular question. However, it is plausible that different sets of rules could give rise to the same pattern. An ABM can show what a set of rules can lead to, but it cannot determine that there cannot be other sets of rules that could do the same. The range of theories that can be used to inform the rule set of an ABM is vast, but a large majority of models rely on some form of bounded rationality (Epstein, 2006). Bounded rationality has the advantage over traditional rational choice assumptions in that it can relax some of the assumptions made about agent rationality. By definition, bounded rationality relaxes traditional rational choice, but has the drawback that the rationality increases in complexity. The strength of Multi-Agent Systems is its ability to model heterogeneous agents with limited capabilities. The downside of this practice is that bounded rationality is not well-defined. Those bounds can be restrictive or relaxed; and are likely to vary widely across models and disciplines. The inconsistencies across disciplines are indeed a concern that needs to be addressed. Several ABM standards have been suggested, but not implemented (Epstein, 2006). A growing interdisciplinary community and a possible conception of a single complexity science field improves the odds of a generalised standards. Until then, we can reduce inconsistency by building on existing models, replicating findings and drawing analogies to sister sciences.

The problem of 'boundless bounded rationality' is not necessarily a problem: firstly, we can repeat the virtual experiments of agent-based models and test the scope of bounded rationality. What may be the source of initial uncertainty can instead be used to generate more knowledge and further our understanding of rationality: what limits there are to its bounds, and how agent behaviour may be influenced by levels of boundedness. Secondly, the vast amount of theoretical and empirical literature on bounded rationality in psychology provides a good foundation from which to make informed choices for models. Herbert Simon (1955) lamented the lack of empirical knowledge in psychology to inform theory. Several decades later we have this pool of resources to put to use.

Computers have advanced, but are still subject to limitations that can frustrate an agent-based modelling effort. Because computers operate sequentially, agents cannot act simultaneously (Gilbert, 2007). A real world experiment may ask everyone in the room to take one step forward; in a computer model, only one agent can act at one point in time (asynchronous execution). This can impact the model outcomes: in the prisoner's dilemma game, synchronous versus asynchronous execution produce a different pattern (Mitchell, 2009). This technical limitation is usually circumvented by placing all agents in some form of a queue, and at every time step, everyone in the queue acts. Because in a queue, the order is predetermined and thus might affect the subsequent decision of agents further down the queue, the queue is shuffled at every time step, so that tactical advantages are kept at a minimum. Time is commonly modelled in discrete time steps, like rounds or taking turns (although continuous time models exist). The relative length of each step can vary depending on the model: in the case of a trading simulation between stock agents, time only progresses whenever something relevant happens. The actual time that may pass between such events in the real world is not relevant. In a simulation of the spread of ancient societies, time steps may constitute years, in each of which the societies can perform some operations. The model description should therefore involve some discussion about the time relative to other parameters in the model.

(iii) Principles of individualism

Many of the aforementioned issues can be traced back to complex nature of the object of study: complex systems are multi-layered, exhibit macro patterns that cannot be reduced to individual components, and feature highly endogenous interactions of the moving parts. Raising these issues is important for good practice, but there have been concerns about agent-based modelling (or microsimulation in general) and its inherent *theoretical* biases that are independent from the nature of the subject that they examine. Agent-based models have the tendency to view the social world from an individualist perspective (O'Sullivan and Haklay, 2000). Because ABMs are made up of individual agents, and these individual agents have rules, attributes and decision-making capacity, ABMs inherently accept the view that society is the aggregation of individual activity. Structural influences and collective social behaviour do not exist unless generated by individual actions- this assumption fails to acknowledge that social structures can pre-exist and can shape individual behaviour (O'Sullivan and Haklay, 2000). There is some disagreement among scholars on whether the English-language research of sociology in the 20th century is biased towards individualism or structuralism (Sawyer, 2005), but it is important to note that the debate exists and what role agent-based modelling plays in that debate. Structuralist sociologists posit that social systems can be understood on their social, or macro-level without the need to consider the individuals that exist underneath (Sawyer, 2005). Social Networks are an often-used example of such a structural phenomenon. The individual-level characteristics are not necessary to further understanding gained from the study of networks. As previously mentioned, the link between the individual and the group in sociology (as opposed to the natural sciences) is difficult to capture. Agent-based models are not free of this 'theoretical baggage'. Brian Epstein echoes these sentiments and laments that in agent-based modelling, researchers (erroneously) overestimate the effect of individual people on macro-level social properties (Epstein, 2011). Models assume that the social entity is entirely composed of individuals. Social structure could be included by modelling agents that belong to different social classes that follow different rules. This way, the social group affects the individual.

There are two aspects to this debate. The first aspect is the existential debate in sociology and to what extent Western thought of individualism has influenced perception of 'the social'. Because Western society has placed an increasing importance on the value and rights of the individual, whose freedoms should be protected from the group, it may cause sociologists to assume that these features are inherent. The inclusion of non-individual entities in an agent-based model is technically achievable, but only if one assumes that no matter how the group exerts influence (collective rules, spatial concentration), the interaction still takes place within the agent, the individual. This is an interesting issue, but it does not warrant the criticism of agent-based models as such. Every method has an inherent bias towards an explanation of the social. An equation-based model derives its strength from unit homogeneity. An agent-based model conversely derives its strength from breaking the whole down to its component parts (Epstein, 2006).

The second aspect of the criticism concerns our understanding of complex systems in general. Complex systems operate at different levels of interaction which may not warrant the binary micro-macro distinction in the first place (Mitchell, 2009). If so, the method of agent-based modelling cannot solve this discrepancy, and it should not be expected to. As any other method, ABMs are a tool designed for a certain task. That task is to generate the macro-state from the micro-state, bottom-up. Philosophical debates and discourse on micro and macro are a separate endeavour which may inform agent-based modellers, but should not bind them to solve a problem in a different domain. To decide on an ABM as a method of scientific enquiry is to accept its weaknesses as much as its strength.

The history of complexity sciences and agent-based modelling has provided us with a point of view that can tackle complex problems in political science. The following two sections will address the substantive elements of this thesis. The subjects of migration, segregation and social pressures through norms are all highly intertwined with today's migrant societies, and warrant a complexity-based approach.

1.2 Immigrant societies

Migration around the world has increased massively over last decades, affecting many spheres of life such as nations, neighbourhoods, workplaces (Citrin, 2015). The 2016 European refugee crisis has marked the largest mass-scale migration into the EU since 1945 (van Prooijen et al., 2017). Voters in Europe and the US have become increasingly worried about the impact of migration and diversity, and populist authoritarian parties have gained greater support in the past 30 years (Inglehart and Norris, 2017). The election of Donald Trump as US president and the UK's decision to vote to leave the European Union in 2016 have signalled a shift in voter attitudes and the rise of polarised attitudes towards migration (Lambert et al., 2017).

As a result, the Social sciences have paid more attention to those questions and concerns about impact of diversity, multicultural ideologies, immigrant adaptation (Ramos et al., 2016). Concerns about migration are not new, however. Research on populism, political correctness and challenging the liberal world order reach back at least two decades (see for example Van Boven (2000), Miller (2000) Foner and Alba (2008)). Migration is not always automatically a problem. But how does it turn into one, or is perceived or framed as such? The challenges of migration are economic, political and social. Economic concerns include costs and benefits of migrants creating and taking jobs, paying for and using healthcare systems, labour market demands and pressures on housing markets. Social concerns include the integration of migrants into the host society, the cultural, ethnic and religious differences of migrant minorities and a change of status quo of previously largely homogeneous societies diversifying (van Prooijen et al., 2017). The different population groups may have no prior knowledge about each other, and the perceived cultural distance may be large. Population groups can be socially excluded and perpetuate segregation of different groups (Arbaci, 2007). The media narrative and political discourse are also part of the cause and effect of immigration problems. Prejudices can shape the way different migrants are perceived and represented in the media (KhosraviNik, 2010). Discourse studies have repeatedly shown how migrants, asylum seekers and refugees are oftentimes conflated and little effort is made in the public discourse to emphasize that refugees and migrants are very different groups with different prerequisites and subsequent demands on the welfare state (Abid et al., 2017).

One focus of this thesis is to draw attention towards the host society and migrants within the host society. Integration is a vital part of migration policy. The focus of integration studies are the migrant and host communities of host countries. What are the challenges of integrating people that come from different cultural and ethnic backgrounds, and how can these challenges be overcome? The structures of immigrant-receiving societies in Europe and elsewhere are changing (Penninx, 2006a). The overarching research question of this thesis is: how does migration affect the societies that are on the receiving end?

In light of this broad question, I want to focus on two related issues. The first is the physical separation of natives and migrants. Does sudden migration act as a shock to the system, causing more segregation of groups? Can different arrival times and pattern reduce shock impacts? Can tolerance of the ‘other’ grow without state intervention? What circumstances lead to higher or lower levels of tolerance?

The tolerance dimension links to the second issue: social norms of tolerance in host societies. With norms promoting multiculturalism, how can public opinion shift so dramatically against it?

Because I model the immigrant society, I draw from the findings of existing literatures to inform my parameter and model design decisions. In the following two sections I review relevant parts of the literature on migration and integration, and norms and attitudes that will form the basis for the agent-based model.

1.2.1 Migration and Integration

Migration is studied in disciplines such as economics, demographics, political science, sociology and anthropology. International migration of human beings traditionally involves the travel from one country (country of origin) to another country (destination country or host country). Migration can be understood as a macro phenomenon of large-scale population shifts and economic costs and benefits; and it can be understood as a micro phenomenon of individual decisions that lead to migration and the challenges that must be overcome to achieve it.

The definition of a migrant has changed over time and can differ from country to country, but in general a migrant is a person that settles down in the destination country; usually for a minimum period of one year (Baganha et al., 2006). The economic study of migration includes the fragmented labour theory of migration: exploring why migrants move; what their economic incentives are and crucially, what the economic gain is for the host country. Because many of the research institutions and grants have come from destination countries, much research is focused on the impact of immigration on these destination countries (Black et al., 2006). Although the volume of research on countries of origin of migrants is smaller, it too has developed and diversified in recent decades. Previous theories had relied heavily on assumptions such as migration being driven by poverty in developing countries; and that low development is the key to understanding high emigration (Black et al., 2006). Whilst poverty can incentivise people to leave their country, it can also provide an obstacle to doing just that: poor people might not have the resources required to migrate (Black et al., 2006). A reduction in poverty can thus facilitate migration, not reduce it. Other studies focus on the migrants themselves as individuals or families of individuals. What drives people to migrate, what determines their decisions of where to go, and what are the facilitators and obstacles in their quest?

Migration is thus a vast subject that has always been explored in some form, but the recent rise of populism, feeding on fears of overpopulation, the notion that the 'boat is full' and that migrants overwhelm the existing population, have highlighted the need to grasp the effect on the host countries. A crucial element of new research endeavours is the perception that migration is here to stay. In order to sustain the growth and economic model of the West, migration is needed to combat the falling birth rates and ageing populations. Minorities are increasing in size; projections suggest that by 2043, White Americans will be outnumbered by the total of non-white minorities (Rios and Wynn, 2016). Migration is thus not just an addition of new, different population groups to existing societies. Centuries-old population dynamics are shifting, bringing change that challenges prevailing conceptions of national identities (van Prooijen et al., 2017), the role of ethnicity, culture, religion in liberal states and their very identity (Joppke and Torpey, 2013). This was not always the case, not even as recently as the 20th century.

Migration in the 20th and 21st century

For many Western (European) countries, international migration has become a major issue in the second half of the 20th century and beyond (Citrin, 2015). Between

1985 and 2005, the number of migrants in the world had doubled. This effect was more pronounced in Europe compared to the US (Novotny and Hasman, 2015). Economically, migration contributed heavily to growth and became an important source of labour for European industry and commerce. The causes, source and impact of migration has undergone many changes throughout the century and consists of many complex relationships: shortly after the second world war, migration to Europe constituted many low-skilled labourers that filled factories and industrial plants. As the century progressed and European economies moved towards increased automation and a knowledge-based society, demand for low-skilled labour decreased. When unemployment in the 1970s rose, the flux of low-skilled migrants became a political problem (Hatton, 2005).

Unlike Canada, Australia or the United States, European countries did not see themselves as immigration countries (Adam and Moodley, 2014). When migrants were admitted to destination countries in the 1950s and beyond, states operated under the assumption that these 'guest workers' would leave once their contracts had been fulfilled[1] (Baganha et al., 2006). When it became apparent that many of these guests had put down roots in their hosts' countries, started families and had no intention of returning, the need for integration policies became an issue: many migrant communities in host countries had formed concentrated communities, or diasporas, that in some cases showed signs of social and cultural exclusion. The consensus is that such exclusive behaviour is undesirable and that migrants should make an effort in integrating in the host society (see Baganha et al. (2006), Hainmueller and Hopkins (2014)). The integration of a large number of foreign-born migrants brought with it a challenge to national identities: 'Who are we?'. These challenges are framed differently depending on the political landscape of each country (Inglehart and Norris, 2017).

[1]In post-colonial countries such as France and the UK, migration from ex-colonies was not assumed to be temporary, but it too was curbed.

Many countries have since shifted their focus on the need for integration. Since the policies of multiculturalism in the 1990s, many European states have moved to the other end of the pendulum, criticising the failure of multiculturalism and whether it is compatible with integration at all (Citrin, 2015). Concerns about migration are part of a wider cultural backlash in the wake of declining existential security (Inglehart and Norris, 2017). The Netherlands and Denmark require newcomers to partake in integration courses if they fail language tests (Baganha et al., 2006). An often used example of a multicultural society with high levels of integration is Canada. The northern neighbour of the US receives mainly highly-educated and well-off migrants. In Europe, many migrants have poor backgrounds, have received little education and often times have very different cultural backgrounds. Of all Western democracies, Canada is the only one in which a slight majority of citizens considers immigrants an asset (Adam and Moodley, 2014).

The increasing diversity of societies had generated renewed interest in the social sciences in the late 1990s. One scholar, Robert Putnam, published a large-scale US study of the impact of diversity on social cohesion (Putnam, 2000). He challenged the prevailing view that diversity is a force for good for social capital and demonstrated that in the US, increased diversity in neighbourhoods was related to lower levels of social cohesion and social capital. His work is credited with spawning a host of new approaches to test his assertions. Researchers from various social science backgrounds would now turn to individual-level explanations of these group-level findings. The field of social psychology had developed theories and findings for decades prior, with an emphasis on intergroup relationships: the study of how different groups of people, however defined, interact.

Integration and intergroup relations

With the advent of diversity and integration-of-diversity studies, a whole social psychology discipline revolving around group identities, prejudices, contact between different groups and threat perceptions has enriched the study of migration. Integration research shifts the focus from economic costs and benefits to the social dimension. These questions touch upon issues such as multiculturalism and have undergone several developments. Scholars such as Putnam (2000) have challenged the view that immigration-induced diversity is a force for good in any and all circumstances. Whilst many of the economic benefits are visible and measurable on a macro level, they do not capture the uneven nature of migration. Migration varies widely in space and time: not all countries receive the same level of migrants, and not all areas in a country receive an even amount of migrants- even within cities, migration occurs in certain clusters (Penninx, 2006b). Temporal shifts might mean a sudden increase of migration in a certain area of a country or city, even though overall migration in that same country has been steady. Whilst migration may be overall a good influence on society, its benefits are not necessarily evenly spread and some people may not feel any benefit at all (Hatton, 2005). Generally, larger shares of immigrants of the population in a country correlate with more negative views on immigration by the native-born public (Citrin, 2015).

Intergroup relations between 'natives' (those born in the host country) and migrants depend on resource availability, public perceptions, cultural distance and potential to 'fit in' (Hainmueller and Hopkins, 2014). Ethnic and cultural differences of migrants influence negative perceptions within the host society (Joppke and Torpey, 2013). In Europe, the most pronounced difference between natives and migrants is ethnicity and religion: many migrants are Muslims and of non-European ethnic background (Joppke and Torpey, 2013). Religion and culture can explain the sharp differences in anti-immigrant attitudes in the US and Europe. Europe's migrants are largely Muslims, the United States' migrants are largely Mexican and thus, largely Christian (Joppke and Torpey, 2013). This reduces the cultural distance to the Christian-dominant US. Perceptions of 'otherness' of migrants change over time. Spanish and Italian guest workers (who were usually white and Christian) were perceived as too different; today, these differences are no longer salient. Before 9/11, Muslim immigrants were Moroccans or Algerians in France, Arabs in the US and Turkish people in Germany. After 9/11, they were all 'Muslims': alien and threatening (Joppke and Torpey, 2013). This shift in group perception is important in understanding why large proportions of migrants have come to be seen as a threat to national identity and culture to many voters. The bounds of what constitutes the out-group (the 'other') can change. Changes of such boundaries can give rise to contesting ways of framing them (Brennan et al., 2013). Populist parties for example can frame these differences based on perceived threats that majority population members can experience (Rios and Wynn, 2016).

Outgroup prejudice towards minorities can be a response to perceived threat to the ingroup (Esses et al., 2008). Contrary to racism, xenophobia has a rational element: when resources are scarce, migrants do compete with locals, and the competition is real (Adam and Moodley, 2014). Thus, despite economic arguments for migration, the notion prevails that the "boat is full" (Adam and Moodley, 2014, p.120). The perception of alien-ness of migrants is reinforced by the large share of Muslims: immigrants are 'doubly different'. The ingroups (non-Muslims, non-migrants) have become more pronounced. Inglehart and Norris (2017) note that the increase of value-based identities over class-based identities has furthermore contributed to a heightened perception of value-based identity.

The range of empirical studies testing the impact of immigration, segregation, ethnic diversity, social capital or intergroup threat is vast, spanning many disciplines. In Political Science, social capital and diversity is a combination that is often used to understand the impact of introducing diversity in a homogeneous society, in part driven by Putnam (2000). Empirical studies have enhanced our understanding of how diversity can impact a society that was previously not accustomed to a diverse crowd. But the jury on explanations of what drives the impact (i.e. the mechanisms) of diversity is still out (Collier, 2013), in part due to the large differences between datasets (Ariely, 2014). Agent-based modelling can be a useful tool to test social theories directly, circumventing the difficulties of data collection and estimation. By modelling mechanisms suggested in social theory and implementing them on an individual-level, an ABM can function as a virtual experiment of the theory and the mechanisms that it puts forward as explanations. A particularly difficult concept for empirical scholars has been social cohesion. It is defined as the 'glue' that keeps a group of people, such as a neighbourhood, together. Cohesion is hard to measure in part because it is constituted of social norms: the unwritten rules that guide the everyday interactions of people. When norms among a group are positive and promote interaction and reciprocity, social cohesion is high (Van Assche et al., 2018). Thus, scholarship that seeks to determine the impact of diversity on the social space of people is bound to study the social norms that provide the context in which diversity increases.

1.2.2 Social norms: a brief primer

The relevance of social norms in the context of our migration society is two-fold: firstly, different norms of different groups set those groups further apart. Different religions promote different set of norms. Secondly, norms regulate how outgroups should be treated. For example, many European countries have social norms that promote tolerance of minorities, including migrants (Blinder et al., 2013). These norms can clash with concerns over cultural incompatibility of migrants.

Different cultures and religions can lead to social friction when members of those groups meet. Visual clues such as skin colour can be easily decoded as ingroup or outgroup attributes. Groups also exhibit 'cultural inertia': inertia contained in beliefs and values influences the decisions that people make (Chong, 2000). Cultural inertia and habits are carried on through the unwritten rules of society: social norms and customs. Some of the differences that distinguishes migrants from the local population of the host society are obstacles to cooperation: if a migrant does not speak the local language, contact with locals will inevitably be more difficult, and contact to other migrants speaking the same native language becomes more attractive.

Norms regulate everyday behaviour without the need for cumbersome laws and can provide a useful heuristic for solving coordination and cooperation problems without having to engage in high-level critical thinking. Some norms solve coordination problems and are beneficial for many participants: by driving on one side of the road, everyone ensures that collisions occur at a greatly reduced rate (Brennan et al., 2013). Norms decrease transaction costs within a social group because they coordinate expectations and choices (Chong, 2000).

Social norms are pervasive because they affect the decision making of people: people will consider what the perceived status quo is given some behaviour (Fields and Schuman, 1976). In survey methodology, the tendency for survey respondents to answer a question in such a way that conforms to an existing idea of what constitutes a morally defensible behaviour or attitude is known as the social acquiescence bias (Krosnick, 1999). Conformity to a group has many benefits: in the absence of information, following the herd is a safer bet in uncertain situations. The more crowded restaurant is more appealing and the empty restaurant, because the presence of people indicates that it is a good place to eat. Heuristics such as these are not accurate, but they are easily deployed and require little effort (Kahneman, 2011). Conformity behaviour is not restricted to solving simple cooperation or coordination problems. In order to ‘fit in’, people will silence themselves (Sunstein, 2003).

Norm conformity is the act of following a social norm- say, offering the seat on the bus to an elderly person because the social norm stipulates that able-bodied should help the less able-bodied so as to relieve their discomfort. Norm defection is the act of defying the norm: remaining seated when an elderly passenger nearby has to stand. Because norm defection carries social risks (stares on the bus, perhaps even apprehension by a fellow passenger), the costs of conformity are usually deemed worth paying (Brennan et al., 2013). What makes social norms particularly persuasive compared to rules of coordination (everyone drives on the left-hand side of the road, thus reducing the risk of collision), is that they shape the emotional response to behaviour or attitudes (Brennan et al., 2013). Knowing that one should give up the seat for an elderly passenger, people will feel bad about defecting. The possibility of punishment is independent from the emotional response. Social norms can be deeply ingrained in our thinking and our moral standards: what is the right thing to think, to say, to do. Norms permeate all levels of society; from holding the fork in the left hand at the dinner table to the justification for declaring war on another country. Norms make acts meaningful: hoisting a white flag on the battlefield is a recognized sign for yielding in battle. Without the recognition of this norm, the act of hoisting the white flag is reduced to the physical act of offering a piece of cloth (Brennan et al., 2013). When a migrant minority does not adhere to the same norms that permeate the host society, this can be interpreted as norm defection.

The tendency to conform and to reduce distance between oneself and other in-group members can result in people adopting a view publicly, although they may disagree with it privately. Once public opinion shifts, non-conformers can pretend to conform. This tendency to adopt the majority view can prevent social change in providing resistance against challenging the status quo. Minority groups can also use this process to push through social changes that are widely rejected. Because people have very little information on what the private attitudes of others are, they guess. And the result can be widespread belief in the dominance of a norm when, in fact, it isn't widely supported. This state in which the majority erroneously believes in the dominance of any given norm, but actually, it is only privately supported by a minority, is called *pluralistic ignorance*. To shift public opinion, periods of pluralistic ignorance are required. Pluralistic Ignorance rose to prominence in studies of racial prejudices in the US (O'Gorman and Garry, 1976). If a majority genuinely believes in white supremacy, people can rationalise their views using the existing norm. Only by misjudging public support for a rejection of white supremacy will adhering to the supremacist view result in a discrepancy with the (assumed) view and thus create social pressure to, if not genuinely swayed, pretend to reject white supremacy, too. The shift away from the political centre towards its periphery and populist parties suggest that this 'sudden shift' was not so sudden after all: attitudes had not been reflected accurately in the public discourse, media attention and to some extent, scholarly research.

1.2.3 Norms of tolerance: not all is what it appears to be

In contemporary Europe, social norms exist that condemn prejudice against specific groups (such as migrants) (see Blinder et al. (2013) and Sniderman and Hagendoorn (2007)). Immigration and the integration of migrant minorities have become salient and contentious issues in many parts of Europe, and anxieties about unsustainable migration levels are widespread (Blinder et al., 2013). In recent years, some populist parties have been successful in framing the migration debate and win the support of large sections of the population (Arzheimer, 2009). Populist parties that are openly discriminating against minority groups are less successful than those that remain ambiguous in their message (Blinder et al., 2013). By remaining ambiguous, these parties can draw the attention of voters which harbour prejudices against certain minority groups, but who are also aware of the dominant norm of non-prejudice.

This is an example of a social norm affecting political choices. Without the norm not to harbour prejudices, support for the restriction of rights of minorities might be more openly expressed. What constitutes a legitimate case for 'discrimination' (for the moment narrowly defined as a legal status that entails fewer rights compared to those of other groups) can vary, too. In some cases, the non-discrimination norm might not be viewed as applicable. For example, citizens of a country enjoy benefits that foreigners don't enjoy. This is seen as a legitimate practice of citizenship, not discrimination (Miller, 2000). Thus, a political party advocating favourable legislation for citizens may not face clashes with an anti-discrimination norm. The process of weighing a potentially norm-adverse position against what is deemed the correct opinion involves estimating of majority opinions. Is non-prejudice supported by a majority? Why is the norm against prejudice *perceived* to be the majority norm? This perception is based on a variety of social clues such as media coverage, public debate, private debate and long-term factors such as education and upbringing (O'Gorman, 1986).

Perceptions of what is deemed acceptable and what is not influence decision making (Sunstein, 2003). If a society is publicly committed to non-discrimination but privately, a large proportion of the populace harbours doubts about this norm, the group can be susceptible to public opinion shifts that are instigated by norm entrepreneurs or otherwise visible actors that can lend a voice to those who have previously silenced themselves (see Wang et al. (2013) for a discussion on the causes of flips in attitudes). The year 2016 was marked by such political shifts away from existing norms of liberalism.

Leave voters in the UK voiced their concerns about immigration and rejected the notion that harbouring doubts about the benefits of migration would brand them racist, or xenophobic. In the United States, Trump voters were not deterred by the number of political gaffes and controversies that their candidate was involved in. 'Political correctness' (PC) as a term has resurfaced in the 1990's as a result of the ongoing battle as to what ideals (Western) society should strive for (Van Boven, 2000). People fight over the right to frame the situation (Brennan et al., 2013). On the surface, liberals had won that battle. The populist surge in Western countries suggests that public opinion had not been swayed as much as it might have previously been believed. In 2000 (!), Leaf Van Boven writes:

> *"The pressure to appear politically correct can have important consequences for social life. In particular, the desire to appear politically correct, and to avoid being seen as racist, sexist or culturally insensitive, can lead people to espouse publicly support for politically correct issues, [...] despite privately held doubts"* (Van Boven, 2000, p.267)

Viewed from this end, the increase in hostility towards the out-group of migrants is not as sudden as it appeared through several large and sudden shifts in the political landscape. Instead, the underlying, privately held beliefs were not expressed until now. We understand what happens when pluralistic ignorance occurs and that it emerges, but we don't know how it emerges in the first place and what circumstances facilitate its emergence. In the context of our complex immigration society, the social pressures of public opinion are a vital component that influence preferences for a segregation of different groups.

1.2.4 Research Questions

The preceding literature review has revealed a number of interesting challenges facing migrant societies. The complexities of migration and integration are enormous and empirical studies have experienced frustration in some areas. The agent-based approach of this thesis seeks to contribute to the understanding of the challenges and questions around migration and its host societies.

In the context of migration societies, there are two research questions that I address, focusing on two different perspectives of the same problem. The first perspective is the physical integration of new people into an existing population. When a status quo is upset, it may take time to reach equilibrium again (or not at all). Using Agent-based modelling, we can look at a hypothetical situation of migration shock-waves that don't correspond to real migration levels, but perception levels. If people act as if migrants were swamping the country, we can model that. Generally, people overestimate the share of the foreign-born or immigrant population (Markaki and Longhi, 2012).

The two broad questions are (i) how does migration affect the host society and migrant population? and (ii) how do norms and social pressure affect attitudes of migrants and out-groups more generally?. The first question is broken down into two specific aspects of migration: the number of migrants arriving and at what rate they arrive; and the initial settlement locations when migrants arrive.

(i) How does migration affect the host society and migrant population?

> Does rate of change of migration affect happiness and segregation?
>
> Does the placement of migrants at the time of arrival matter?

(ii) How does norm conformity to be tolerant of minorities interact with growing discontent with the status quo?

> What gives rise to pluralistic ignorance in this context?

1.3 In this thesis

The previous section has reviewed the existing literature on migration studies and social norms. Migration, in itself a complex system, affects existing societies in many different ways. This makes the migrant society a very good subject to approach with agent-based modelling. We can construct our complex immigration society in which agents are exposed to migration, must cope with outgroup contact situations and navigate through the social space whilst conforming to the existing norms. The importance of private beliefs and public pressures in an immigration society concludes the substantive interests of this thesis.

I show that at different levels of analysis, agent-based models allow us to understand possible mechanics behind the macro-patterns that we observe in the real world.

The baseline agent-based model that will be used is Schelling's model of self-sorting behaviour (described in greater detail in Chapter 2). Using this well-established and widely applied model improves our ability to compare generated results and provides us with a breath of robustness checks that have been applied in other versions of the model.

The first subject of interest is segregation outcomes under conditions of migration. Chapter 2 introduces the migration of new agents onto an existing grid, modelling migration into a social space, which is a novel addition to the literature. In this version, segregation and overall population happiness are the primary outcomes of interest. This chapter demonstrates the depth of findings that can result from a simplistic and idealized agent-based model that trades theoretical accuracy for parsimony. The theoretical basis for the model design is drawn from a range of existing findings from the political science literature.

Chapter 3 builds on this migration concept and moves towards the inclusion of literature on norm conflicts that can exist between hosts and migrants. The model in this chapter includes the crucial addition of adaptive tolerance, an addition designed to implement the contact theory into the model. Agents are heterogeneous and part of a tolerance feedback loop. Segregation preferences are no longer predetermined. Preference development is now based on contact theory. This model version demonstrates the ability of agent-based models to test existing theories by implementing them in a repeated virtual experiment.

Chapter 4 moves on from the migration aspects of the previous models and focuses on the social norms that can pressure people to adapt their preferences, at least outwardly. This model represents a smaller-scale social space in which movement is severely limited and existing social structures exhibit great influence on agents. This adaptation of the Schelling model highlights the ability to model latent behaviours and their long-term effects, and how these models can be used to enhance a theoretical debate. The subject of interest in this chapter is the emergence of pluralistic ignorance. The study of emergence, central to complexity science and subsequently, agent-based modelling, can contribute to our understanding of political behaviour.

2 Chapter 2

In this chapter I introduce a version of the Schelling model of segregation that includes a mechanism of immigration. The Schelling model is part of the "contextual neighbourhood effects" approach to diversity and segregation which posits that the context of the neighbourhoods in which the variables of diversity, social cohesion and segregation are measured, has to be explored rather applying existing theory that might omit these effects (Andersson et al., 2017). Because the impact of diversity is not straight-forward, the boundary conditions have to be explored first (Ramos et al., 2016).

The model simulates a host society of existing 'native' agents which is exposed to 'migrant' agents entering the existing society. The chapter addresses the first set of research questions: How does migration affect the host society and migrant population? Does the rate of change of migration affect happiness and segregation differently, and does the placement of migrants at the time of arrival matter?

The Schelling model in its current form was published in 1971 by Thomas Schelling in a bid to understand how individual decisions of agents to relocate could lead to a macro pattern of segregation (Schelling, 1971). In the context of widespread racial segregation in the US, the model could demonstrate that for segregation to occur on a macro scale, no deeply entrenched racism was required- even slight preferences to reside with people of ones' own colour could lead to segregated areas. The model has since been adapted and advanced in multiple ways and is a well-known model of self-sorting behaviour. A recent adaptation by Hatna and Benenson (Hatna and Benenson, 2015a) incorporates assumptions of a heterogeneous society in which preferences for friendly neighbours would vary. This chapter builds on their model and investigates how far the rules can simulate patterns of migration in segregated cities. The literature on migration has not featured a Schelling model implementing external migration onto the existing grid (rather than already existing agents migrating within the grid).

Migrant waves in the past have been accompanied by public debate in the respective host countries (Baganha et al., 2006), and there is no consensus as to whether immigration into the host societies is, as a whole, a positive or negative influence (Penninx, 2006b). Because the subject is vast and complex, the academic debate has diverged into the respective disciplines: the economic impact of migration can differ from the sociological or ecological impact. In the past decades the notion that immigration is generally good for the host population has been challenged (Putnam, 2007). Subsequent analyses have drawn heavily from pre-existing sociological and socio-psychological studies on group identity, racial relations and segregation. The latter subject area has been influenced by the Schelling model because it had been so successful in demonstrating that segregating behaviour can result even in a moderate population as a result of moderate preferences to be in the majority group of the population. The theoretical justification of the Schelling model design are found in social psychology rather than sociology, because the individual-level model design lends to individual-level theories of how people (agents) act, react and interact (Ramos et al., 2016). The model design decisions are based on empirical findings and stylistic facts that social science models have developed. This serves as a way to explore the interactions between the new addition of migration with the existing model components.

The migration literature has enjoyed a host of agent-based approaches (see Klabunde and Willekens (2016a) for an overview), but most attempt to explain *why* migration occurs. In the scenarios that this chapter considers, migration is taken as given, but its intensity (how often does it occur, and how many migrants arrive) and the makeup of incoming migrants (where do they arrive) differs. The goal of this chapter is two-fold. Firstly, it seeks to add to the theoretical insights that the Schelling model can give us, not just for migration but for general behaviour under conditions of sudden external shocks. Secondly, it seeks to evaluate whether and how the group of newly arriving agents affects the future pattern of segregation in the population. In order to simplify replication of experiments, the model is based on the aforementioned Hatna and Benenson study. Based on their description of the model, I recreate their model and then proceed to adapt the model in order to address my research questions. This approach, to focus on the rate and change of migration, shifts the focus on the host population.

The structure of this chapter is as follows. The first part briefly summarises previous research done in the areas of migration and ethnic segregation, and how Schelling models in general operate. Subsequent paragraphs discuss how Schelling models and migration can be combined, and how the model implements the migration element. Afterwards, the collection of data from the simulations is discussed and the results are presented through analysis and a brief discussion follows at the end.

2.1 Migration and ethnic segregation

The combination of migration and ethnic (residential) segregation has an intuitive appeal. Immigrants tend to cluster spatially (Bushi, 2014), and so do ethnic groups (Søholt and Lynnebakke, 2015). Migration is usually defined as the movement of people from one place to another. These places can be countries, regions, boroughs, cities or neighbourhoods. The type of migration that is of interest to this chapter is international migration of people from their country of origin to another country (host country). When migrants enter a country, their point of entry is not random. Cities such as London have distinct areas that are well-known for accommodating newly arrived migrants (Hall, 2013). International migration is still increasing (UKCensus, 2011) and affects the ethnic makeup of global cities such as New York or London.

Ethnic segregation occurs when people perceive a group of other people as different based on ethnicity and subsequently seek to live in closer proximity to people more like themselves. Migrants are an obvious group that can be singled out as different since they are foreign to the country. This can be, but should not be, linked to differences in ethnicity. Visually poignant features such as skin colour make it easy for people to distinguish between those alike and those that are different. Studies on migration and ethnic diversity are widespread. Putnam finds that migration can increase the social costs of cooperation if the resulting society is more diverse (Putnam, 2007). The proposed link between high diversity and low social capital has since been tested, yielding contradictory results, most likely due to differences in operationalisation of the social variables (see Ariely (2014) for more discussion). Segregation is to some extent a natural phenomenon: families live together. Members of a family often share the same ethnic group. Second- and third-order relatives may also live in the vicinity, and share the same ethnic background. Such an overlap in cleavages can result in segregation without any explicit preference to *stay away* from a different group (Schelling, 1971). When people segregate along ethnic lines, it often has adverse effects on the community, usually the minority of lower social standing (Zhang and Jager, 2011).

Figure 2.1 shows the different population density of white British, Black/African/Caribbean and Arabic ethnic backgrounds across London. There are visible clusters of different ethnic groups in the different London boroughs. White British are the most numerous, but they are highly concentrated on the outskirts of the city.

Figure 2.1: The ethnic makeup of London in percentages, based on 2011 Census Data. Colours represent population densities of ethnic groups.

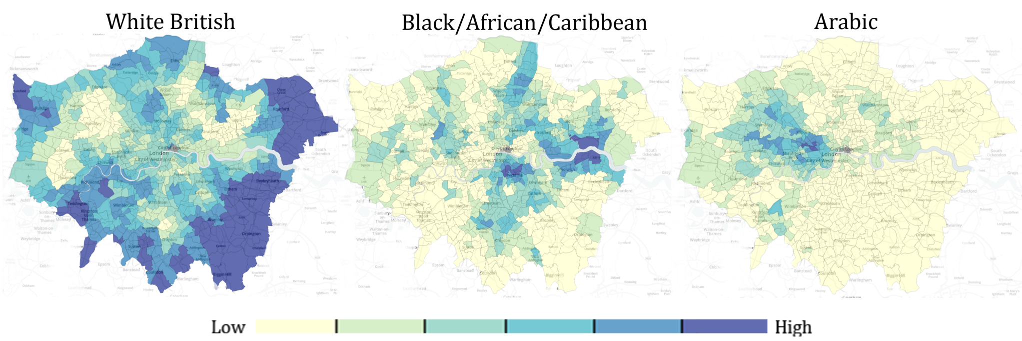

Migration is likely to make a difference to patterns of segregation. In a separate literature that is primarily concerned with international migration, a subject of interest oftentimes are diasporas: pre-existing communities of foreigners that exert a form of attraction to fellow countrymen and women to move to the diaspora (Collier, 2013). Diasporas thus grow larger and faster after forming, until a point is reached at which the host population grows weary of its size and spread, and political measures are employed to reduce the growth of diasporas. Diaspora growth has been linked to the gravity model of migration: migrants are pulled towards already existing migrants, even in the absence of pre-existing family ties (Klabunde and Willekens, 2016a).

It is thus an intuitive conclusion that the rate of flow of migrants is at least in part a function of the existing "map": what country they move to and how the ethnic make-up of a country is shaped. Countries without diasporas are less attractive to migrants (Novotny and Hasman, 2015). Equally, the ethnic makeup of a country or city is (in part) a function of the rate and flow of incoming migrants. There is a vacuum in research on the impact of migration on the happiness on the host population (Collier, 2013). The aim of this chapter is to fill part of this vacuum by exploring a possible mechanism for impacting host societies and immigrant minorities.

The research questions focus on several well-known variables that are used in migration research, such as the rate and change of migration flow, the size of migration communities and the settlement patterns of migrants entering host societies. I break down the broader questions into the following:

1. How does the rate and change of migration impact the happiness and segregation behaviour of the population?

    a. How does the place of arrival of migrants affect long-term segregation and happiness patterns?

b. How does the size of groups in relation to one another interact with migration?

2. Does the placement choice of migrants alter long-term behaviour observed in the overall population?

2.2 Method

In line with one of the two purposes of this thesis, the research questions laid out above will be tackled using agent-based modelling. Following calls for standard model practices (see Collins et al. (2015) and Bruch and Atwell (2015)), I will be using the Schelling model, a well-known model to test segregation behaviors. Before describing my adaptation, I shall review the nature of the Schelling model below.

2.2.1 The Schelling Model

Thomas Schelling published the first version of his model in 1969 and revised it in 1971 (Schelling, 1971), introducing the 2-dimensional model that is well-known today. The model describes a self-sorting mechanism of different groups of people, or agents. In the context of a highly racially segregated United States, the model was able to generate, using very few and simple rules, patterns akin to those of urban segregation.

The basic premise of the Schelling model is that there exist two equally-sized groups of people, usually represented through different colours, in a two-dimensional world. All agents have a common preference as to how many 'non-friends' they tolerate in their vicinity. 'Friends' are agents of their own colour. All agents have the same preference f, the fraction of like-coloured agents that is preferred in their immediate neighbourhood, and the fraction of non-friends is $1 - f$. Say, $f = .5$

denotes that an agent accepts 50% of other-coloured agents in the neighbourhood. Agents cannot see the entire world, only several tiles in each direction. There is no stacking of agents, every tile can either be empty or hold a maximum of one agent. Agents can move around the grid without any restraints. Movement is costless. Agents will move if they find a neighbourhood with more friends than the preference f stipulates. Figure 2.2 shows a visualisation of a Schelling model. In this case agents are blue or green, and white space is empty. The grid is usually at near-capacity of 98% of the grid covered with agents. The usual sequence is that agents are first randomly placed on the grid. Then, each turn, every agent gets to act: they will evaluate whether the current neighbourhood meets their f threshold, and will move when it does not, and stay when it does.

Figure 2.2: Visualisation of a Schelling model: Blue and green agents are first randomly positioned. They then move in accordance with their preference. After a while, clear clusters of segregated groups are visible.

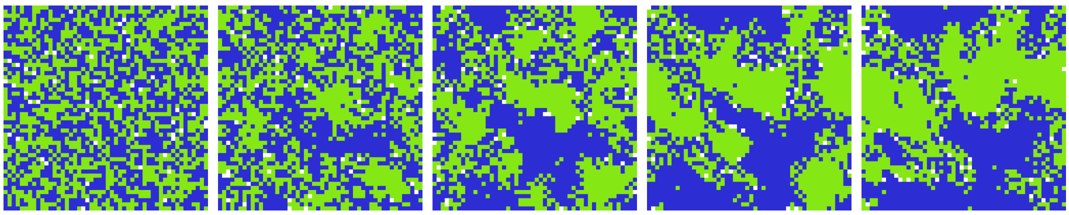

The Schelling model has been studied extensively (see for example Singh et al. (2009), Singh et al. (2011), Shin et al. (2014), Cortez et al. (2015)) and there are a number of things that we know about model outcomes and usual patterns. The model shows critical thresholds at $f = .75$ (Benenson and Hatna, 2011). At .75 (three-quarters of the neighbourhood must be friends), agents segregate into large groups with unhappy agents that reside at the fringe of each cluster. Any f value beyond .75 results in a breakdown of segregation and no equilibrium is reached. Agents are always unhappy, as their demands are never met, and will move around continuously. Segregation patterns become visible at $f = .2$ but are still outnumbered by mixed, or integrated, agents. At $f = .3$, agents segregate visibly (Benenson and Hatna, 2011).

Schelling showed that people do not need to be extremist in order to live in segregated neighbourhoods- in fact, even when people prefer mixed neighbourhoods, the aggregate pattern still tends towards segregation. For example, a person that tolerates mixed neighbourhoods, but wishes to be part of the majority (i.e. $f > 0.5$), the aggregate pattern tends towards segregation: if everyone wants to be in the majority, it is not possible to live in a mixed neighbourhood and have all people satisfied. Recent studies interviewing migrants confirm such preferences for migrant communities as well (Søholt and Lynnebakke, 2015).

The model generates *small-worlds* in which patterns of behaviour do not scale with the size of the lattice (Singh et al., 2009). When the grid is increased from 50×50 tiles to 100×100 tiles, the size of the segregated clusters does not increase. Instead, more clusters of the same size observed at 50×50 emerge. This restriction is due to the limited vision that agents have.

Because the Schelling model was able to provide an explanation for the mechanism of persistent racial segregation in the US, the model has enjoyed considerable attention since the 1980s (Clark and Fossett, 2008). The persistence of segregation is the result of tipping behaviour which is driven by the coordination problem of every agent aligning themselves in such a way that their preferences are satisfied. Schelling termed this behaviour 'speculative evacuation' (Schelling, 1971, p. 185): when a predominantly white area receives an influx of blacks, white people might assume that an eventual tipping of balance towards a black majority is inevitable. They vacate the area, thereby lowering the required number of blacks needed to tip the majority balance. This means that "[...] integrated residential patterns are inherently unstable" (Zhang and Jager, 2011, p. 169). Segregation persists because it is a stable state: either majority is the favourable outcome for each respective group, and the dynamics of the behaviour are such that the majority will always be sought. Integrated residential patterns on the other hand are inherently unstable (Zhang and Jager, 2011). The ease with which segregation can arise and sustain itself is important in our migrant society. If speculative evacuation is the driver for segregation, it may not make a difference how migrants have entered the existing society; how many they are and how quickly they got there.

The Schelling version I use is based on a recent adaptation of the model which has successfully generated both segregated and integrated areas. Hatna and Benenson (Hatna and Benenson, 2015a) introduce a heterogeneous society in which agents can have one of two different preferences. This is more realistic than Schelling's assumption that everyone has the same minimum threshold. They are able to show that with two preference groups, the setup can generate patterns of both integrated and segregated areas on the grid, which is in line with real-word census data of US cities 2010 that they cite: usually cities consist of both segregated and integrated areas (Hatna and Benenson, 2015a).

The crucial addition to their model that I make is the introduction of migration. Not all agents will be generated at the start of the simulation, but some will enter later on. I vary the rate of migration (how many migrants enter at once, and how often) and the type of arrival (do migrants cluster, or disperse). I will return to ethnic composition and the importance of minority-majority relations later on in this chapter.

It should be noted that Schelling himself discussed the possibility of introducing out-group agents via inward migration (Schelling, 1971, p.161). Testing one setup of migration, Schelling observed that clusters only form once enough out-group agents have arrived. At this point, Schelling moves on to a different treatment and does not pursue the effects of migration further. This chapter will pick up where he left off and test whether and how migration affects the segregating behaviour of agents in the Schelling model.

2.2.2 The model

A set of agents $A_t = \{a_1, \ldots, a_{n,t}\}$ are located on a toroidal grid with a total of $N = 50 \times 50$ tiles at time $t \in \mathbb{Z}$. Each agent a_i has a colour attribute denoted c_i, which is either blue ($c_i = B$), or green ($c_i = G$) and they occupy one tile on the grid, location $p_{i,t}$. Green agents are the hosts ('natives') and they are randomly placed onto the lattice at the beginning of each simulation. Natives will occupy the grid when initialised. The division of the population into two distinct groups serves the purpose of representing in- and out-groups as experienced in the real world and described in political psychology literature. An in-group is the group to which a person feels they belong to, the group they identify with. The markers for identification can be externally determined (such as ethnicity) or can be acquired through active engagement or a shared interest (such as a particular social standing) (Pettigrew et al., 2011).

The $NatShare$ is also the starting density of occupation of the grid. Blue agents are migrants and arrive at a later stage. Migrants will continue to arrive until the target density $FinalDen$ is reached and both groups make up half (or a proportion otherwise specified by $NatShare$, described below) of the total population. Agents cannot die or otherwise exit the grid. Population density is a crucial element to the model, as it serves as a proxy for real world choices in the housing market and in urban areas in general. Different type of housing is attractive to different kinds of people. The availability of different housing types can drive segregation in urban spaces (Skifter Andersen et al., 2016). Thus, the Schelling model is usually operated using very high densities around 98% to approximate the reality that people cannot simply pick a place and choose to live there, but rather are constrained by housing availability and their personal wealth to meet the costs associated with moving.

Moving is free in this model, but agents have a preference to be near other agents of the same colour. This serves as a way to simulate urban behaviour and to avoid unrealistic spreading of all agents into the periphery so as to avoid any friction. Agents can only see their local neighbourhood. The model uses a Moore neighbourhood[2] of 5×5 tiles. Each agent's neighbourhood consists of 24 tiles. Formally, let $N_t(a_i)$ denote the neighbourhood of a_i at time t, which consists of the set of all other agents located on the lattice within a Euclidean distance of two nodes from a_i (allowing for diagonal movement); that is, the neighbourhood for a given agent consists of all other agents within the 24 nearest locations to it. This definition of the agents' neighbourhoods is derived from the baseline model by Hatna and Benenson (2015b) and was chosen to ease the comparison to other similar models. Different neighbourhood sizes change the ratio from neighbourhood to the total grid size, thus influencing the freedom of choice that agents have.

Each agent a_i has a tolerance threshold $F_{i,t}$ which determines the number of in-group members the agent desires in their immediate neighbourhood. Higher values of F therefore correspond to lower values of tolerance. An agent with $F = 24$ will not tolerate any out-group member in the neighbourhood. By default, each population group consists of half $F_1 \sim U(0, 7)$ and half $F_2 \sim U(17, 22)$ agents, where $\sim U(17, 22)$ denotes the uniform distribution over the interval 17, 22. The differing ranges of F ensure that one group is tolerant below the critical value 12 that leads to segregation, and on or above the critical value 17 that prevent equilibria due to the inability to find satisfactory locations (Hatna and Benenson, 2015b). I use the F(riend) notation to replicate Hatna & Benenson's notation and to emphasise that the range intervals are based on their findings of critical values. The range of 0-24 refers to the number of potential neighbours within the fixed neighbourhoods of 24 tiles.

[2]Moore neighbourhoods are square neighbourhoods of $n \times n$ tiles where $n \geq 3$.

Agents are either satisfied or dissatisfied with their neighbours. They are satisfied if and only if the number of nearby similar agents, $s_{i,t}$ is at least as great as their desired number of friends. The utility of agent i at time t is denoted $u_{i,t} \in \{0,1\}$ and is given by:

$$u_{i,t} = \begin{cases} 1 & : \quad s_{i,t} \geq F_{i,t} \\ 0 & : \quad s_{i,t} < F_{i,t} \end{cases} . \tag{2.1}$$

Algorithm 2.1 Movement rule for agent a_i

$L \leftarrow$ RandomlyChooseVacantSites(z) ▷ choose $|L| = w$ candidate locations
$L\star \leftarrow \{p_{i,t}\}$ ▷ initialise satisfactory locations
for all $l \in L$ **do**
  $s \leftarrow |\{a_j \in N_t(l) : c_i = c_j\}|$ ▷ number of friends in neighbourhood
  $d \leftarrow |\{a_j \in N_t(l) : c_i \neq c_j\}|$ ▷ number of non-friends in neighbourhood
  if $s_{i,t} \geq F_{i,t}$ **then**
    $u_{i,t} = 1$
    $L\star \leftarrow L\star \cup \{l\}$ ▷ update satisfactory locations
  end if
end for
$l\star \leftarrow$ ChooseOneAtRandom($L*$)
$p_{i,t+1} = l\star$ ▷ update location

The set of dissatisfied agents is given by $D_t = \{a_i \in A_t : u_{i,t} = 0\}$. At every time period each dissatisfied agent $a_i \in D_t$, who is currently located at $p_{i,t}$, randomly samples $z = 30$ of unoccupied locations L_i from the lattice.[3] They then randomly choose a new location from the subset of these for which the number of friends meets their friend threshold $\{l \in L_i \cup p_{i,t} : s_{i,t} \geq F_{i,t}\}$. If no satisfactory alternative locations are found,then the agent remains at its current location $p_{i,t}$. The movement rule for an agent a_i is summarised by the pseudo-code in Algorithm 2.1.

As in Benenson and Hatna (2011), with probability 10^{-2} per tick, agents that are satisfied will also relocate, this time randomly picking a location from z randomly-

[3]The value 30 was adapted from the model by Hatna and Benenson (2015b). The authors have not explained the choice of 30 tiles. I ran N=1,000 simulations testing a range of $z \sim U(25, 1250)$ which did not affect any of the model outcomes. I thus adapted the value of 30 to keep the models consistent.

chosen vacant locations, without considering their utility. This models the fact that people in the real world will move due to a variety of reasons, and not just due to diversity tolerances. Moving is also not purely down to happiness: many life situations call for moving, even if the present situation is satisfactory. If an agent fails to find a new location, the agent will remain at the current location. If the agent is still unhappy in the next time period, it will try and find a new location again.

The concept of moving away when unhappy, unhappiness in turn based on neighbourhood diversity has empirical merit: people's intention to move (either within the city or to another city) depend greatly on the general satisfaction (Van Assche et al., 2018). In real life, this satisfaction is not just influenced by neighbourhood diversity, but also by the social norms that maintain the cohesion in the local community (Van Assche et al., 2018). These norms are not modelled in this version: their absence is interpreted as the absence of positive social norms: because agents will not be swayed to remain in a diverse area, they have no social capital to 'buffer' against the effect of diversity. Equally, agents do not possess any form of ideology, education or economic status that would otherwise influence the perceptions of diversity (Van Assche et al., 2018). As such, the model tests the lower bounds of the parameter range.

2.2.3 Initial Conditions

The major contribution of this chapter is the introduction of migration into the Schelling model. Migration occurs in waves, described in greater detail below. The migrants will settle on empty tiles, and then proceed to follow the same rule-set as the native agents. In order to make way for incoming migrants, the assumption that urban neighbourhoods operate at a near-full capacity (i.e. nearly 100% of the grid is covered with agents) needs to be relaxed. The starting density is therefore lowered. I test the impact of migration waves and the impact of arrival locations in three separate treatment rounds which will be discussed in turn below.

The remaining parameters are listed in the tables below. Note that the values for constants in Table 2.2 refer to the default values when no treatment is applied. Table 2.3 lists the state variables referenced in the model description.

Table 2.1: Independent variables

Parameter	*Distribution*	*Description*
$FluxType$	$\in Scatter, Diaspora, Cluster$	Mode of arrival for migrants
$FinalDen$	$\sim U(75, 98)$	Final density of the population
$NatShare$	$\sim U(2, 98)$	Native share of the population
$F1$	$\sim U(0, 7)$	Desired friends of $FShare$ of each group
$F2$	$\sim U(17, 22)$	Desired friends of $(100 - FShare)$ of each group

Table 2.2: Constants

Constant	*Description*
$t_{max} = 2,000$	Maximum number of ticks per simulation
$N = 50 \times 50$	Size of lattice
$z = 30$	number of considered empty tiles

Table 2.3: State variables

Variable	*Description*
A_t	The population of agents
$c_{i,t}$	Colour of agent a_i
$F_{i,t}$	Tolerance of agent a_i
$u_{i,t}$	Utility of agent a_i
$p_{i,t}$	Position of agent a_i
$N_t(a_i)$	The set of agents that are neighbours of agent a_i
$N(p)$	The set of locations in the neighbourhood of location p

For each simulation, the initial conditions are stipulated by:

1. E The number of waves of migration which specifies how many migration events occur and how many migrants arrive each time. The different waves are described in Section **Waves of migration: E** below.

2. $FluxType$ The settlement behaviour of migrants which denotes Where migrants will locate to at the time of migration. The baseline model uses a

Table 2.4: Dependent variables

Variable	*Description*
M_t^c	Segregation of colour at time t (equation 2.8)
M_t^F	Segregation of tolerance at time t (equation 2.9)
M^c	Segregation of colour at the end of the simulation
M^F	Segregation of tolerance at the end of the simulation
$\bar{h}_t$	Average happiness of agents at time t
$\bar{h}_t^G$	Average happiness of natives at time t
$\bar{h}_t^B$	Average happiness of migrants at time t
$\bar{h}$	Average happiness of agents at the end of the simulation
$\bar{h}^G$	Average happiness of natives at the end of the simulation
$\bar{h}^B$	Average happiness of migrants at the end of the simulation

density-based settlement type: newly arriving migrants will seek the most densely populated areas on the grid, regardless of whether this dense area consists of migrants, natives or both. Migrants will also try and form a coherent unit when arriving together. There are two more types of arrival, which will be outlined in Section **Arrival locations: FluxType**.

3. $FinalDen$ The final population density that is reached once all migrants have arrived. The amount of empty space that agents have in Schelling's model is significant: the more space there is for intolerant agents to free themselves from unsatisfactory neighbourhoods, the more likely they are able to avoid dissatisfaction. In general, more agents can be happy and much more quickly in situations of empty space. Crowded places on the other hands introduce friction, as agents are not able to escape as quickly (if at all).

4. $NatShare$ The share of native agents of the population. High levels of native share translate into a minority situation for migrants. $NatShare$ ranges on the interval $(2, 98)$, testing a range from a small native minority of 2% to vast majorities. Because Schelling models don't scale well beyond the size of cities (Singh et al., 2009), the grid is assumed to be no larger than a city. Country levels of migration are usually low, but cities can experience high migration population share, which can be modelled by decreasing $NatShare$.

Waves of migration: E

There are four different migration treatments and the control, no migration. The number of native agents is always constant, but new migrants can arrive in discrete waves of migration. The treatment is summarised in Table 2.5 below.

Table 2.5: First round treatment conditions

Treatment	$E = 0$	$E = 1$	$E = 4$	$E = 15$	$E = 100$
Migration	No	Yes	Yes	Yes	Yes
Number of waves	-	1	4	15	100

A visualisation of the different migration treatments is shown in Figure 2.3 below. Each row displays a treatment, the control experiment without migration is at the top. The control in Figure 2.3a is closest to the standard Schelling model. Agents of two colours are randomly initialised, and as they start to move, they segregate into clusters of their own. Figure 2.3b places all migrants at once, effectively creating a time-lagged initialisation because natives don't segregate if there are no out-group members around. When the rate of migration increases in Figure 2.3c, 2.3d and 2.3e, the size of each wave shrinks and the grid is sparsely populated for longer periods of time. The higher the rate, the longer will migrants be a minority population.

At the beginning of each simulation, a total of N_G native agents are placed randomly onto the grid, where[4]

$$N_G = \text{round}(NatShare \times N), \tag{2.2}$$

and for the treatment where there is no migration ($E = 0$) a total of N_B migrant agents are also placed randomly, where

$$N_B = \text{round}((1 - NatShare) \times N_{max}) \tag{2.3}$$

In treatments with migration, there are no migrant agents on the lattice at the

[4]The round function rounds to the nearest integer: round up when decimal is $\geq .5$, else round down.

Figure 2.3: Screenshots of each migration treatment and the control.

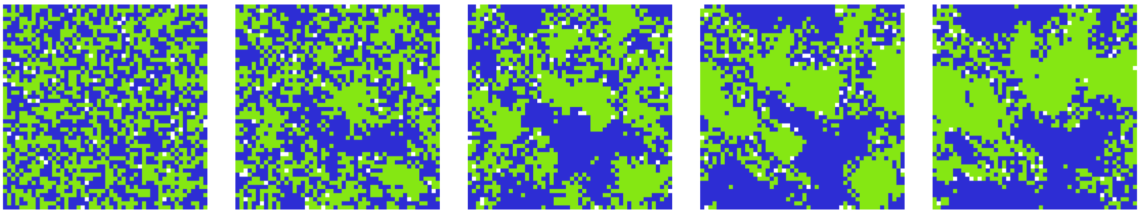

(a) $E = 0$. The map is populated with both groups from the start. No migration.

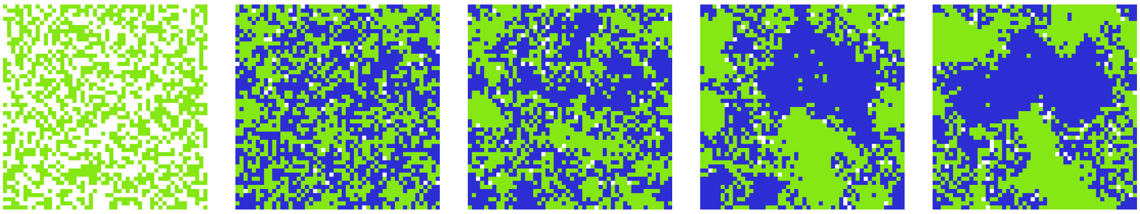

(b) $E = 1$. The map is populated with natives. Migration occurs once.

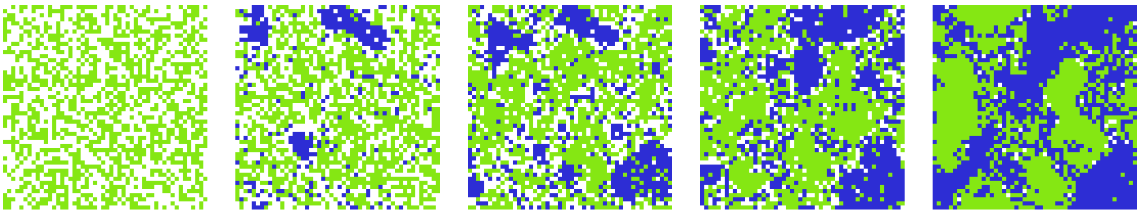

(c) $E = 4$. Migration occurs in 4 instances, with $\sim$306 migrants each time.

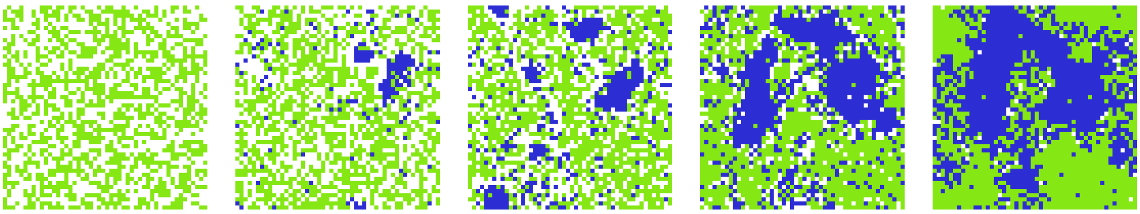

(d) $E = 15$. Migration occurs 15 times, $\sim$82 migrants enter every time.

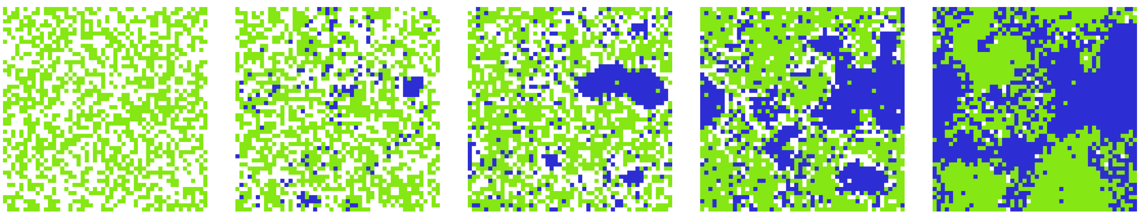

(e) $E = 100$. Migration occurs 100 times, $\sim$12 migrants enter every time.

beginning of the simulation. The first wave of migration occurs at time $0.05 \times t = 1000$, and the subsequent migration waves occur at evenly spaced intervals of $0.9 \times t = 1000/E$ ticks. Migration is restricted to the first 1000 rounds so that long-term behaviour that is not still recovering from a shock can be drawn from the results. During each wave of migration an additional number

$$\Delta_B = \text{round}(N_B/E) \tag{2.4}$$

of migrant agents are placed sequentially (between turns, so that in effect, they arrive simultaneously) onto the lattice, clustering around a focal vacant location p_T. The focal location is used as a starting point, and subsequent migrants will be placed on vacant attractive tiles nearby. What determines the choice of the focal point and the attraction of nearby vacant tiles depends on the $FluxType$, described in Section 2.2.3 below. To demonstrate the basic principle of clustering migrants on arrival, I use the heuristic of clustering around dense areas. With this $FluxType$, referred to as $Cluster$, the objective of incoming migrants is to settle in densely populated areas. Whether this area consists of migrants or natives is not relevant. This is in line with our knowledge that migrants usually target job-rich environments which are more common in cities, moving on to the next-best city if opportunities decline (Collier, 2013). The additional migration locations for the new arrivals are chosen by iteratively finding the best neighbour of the chosen focal location p_T; sites are ranked firstly according to the highest number of surrounding *new* migrants, and secondly according to their local population density within their neighbourhood. That is, the incoming migrants exert an attraction to each other while migration is occurring. This ensures that migrant waves will form clusters whenever possible.

Figure 2.4 shows examples of the influx mechanism at work. This is not the actual model, but a demonstration of what happens during the placement of migrants. The first row 2.4a shows the most basic example: a completely empty map is populated. For illustrative purposes, agent movement is disabled. The cyan-coloured tiles are the tiles that will be populated by migrants in the future, so they are 'flagged' for future arrivals. In the next screenshot the migrants have arrived (blue).

Figure 2.4: Screen captures of the influx mechanism (using density-clustering) in progress. The darker the shade of grey of an empty tile, the higher its appeal rating to migrants.

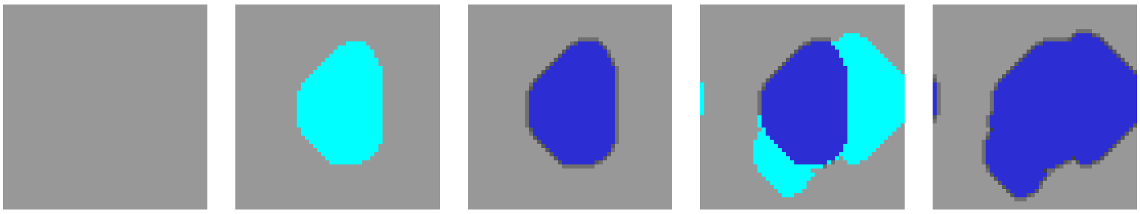

(a) A completely empty map is populated over time. Agents are stationary.

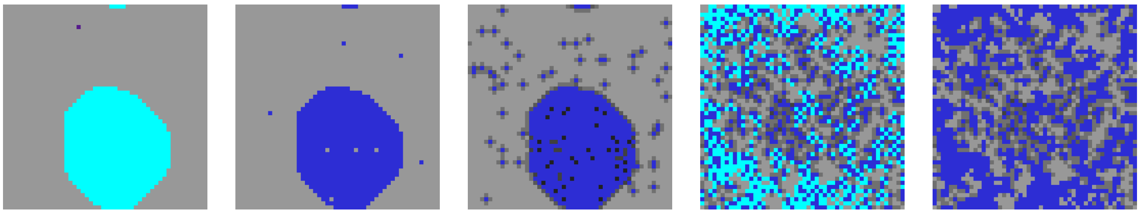

(b) An empty map is populated over time. Agents move around randomly.

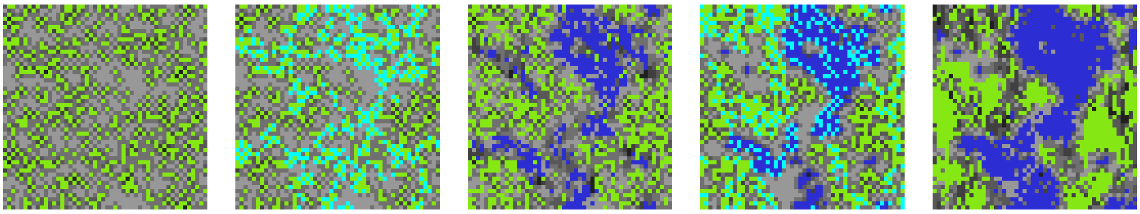

(c) An already-populated (30%) map is populated so that the density reaches 80%.

The existing agents (in this case, just migrants) exert an appeal to future migrants, and thus, their neighbouring tiles are a darker shade of grey. In the following migration wave, the cyan-coloured tiles start off by surrounding existing migrants, but the cyan tiles exert an appeal themselves, so that the future settlement will always prefer a cluster when possible. The cluster of tiles contains random elements by randomly picking a neighbouring tile that is not directly adjacent to the existing cluster. That is, instead of the circle increasing line by line, that line of tiles can sometimes go off in a different direction away from the cluster. This prevents a perfect sphere and provides more natural clustering.

The second row 2.4b shows the same process as above, but this time, agents that have already arrived, have a 1% chance of moving randomly throughout the map. After the first settlement, several migrants have moved from the cluster. Their appeal rating is visible around each of them. By the time the second migration wave occurs, the cluster has completely dissolved and existing agents are randomly scattered throughout the map. The incoming migrants now arrive in smaller, dispersed clusters. The last screenshot in this row shows the dark grey areas in the population centre, revealing its appeal to future migrants.

The third row 2.4c demonstrates the influx mechanism with existing agents, the natives (green-coloured tiles). I recall that for this demonstration, initially, migrants go to populated areas, and it doesn't matter to them whether the population is made up of natives or migrants. However, eventually they move near friends to increase happiness; so natives dont exert the same appeal except initially. The first migration wave will form the same snaky clusters as shown in 2.4b. In this particular instance, both groups favour large majorities for themselves and segregation occurs quickly. The newly arriving migrants are usually placed in the empty space between the two groups, and will quickly move away and increase the size of each respective cluster.

The placement algorithm is summarised in the pseudo-code given by algorithm 2.2. The process of migrating begins by gathering all empty tiles, or locations. For all empty locations, the neighbourhood density is calculated. The highest density neighbourhood will be the starting point for migrants. When more than one *bestStart* is determined, one of them is chosen at random. As long as there are still migrants to place in the current wave, the next-best neighbour from the *bestStart* is chosen using algorithm 2.3. If the neighbourhood is full, tiles from surrounding areas are chosen so that the migrants will always form a cluster when possible. A visualisation of this process can be found at Urselmans (2017b).

Because the number of waves is pre-determined, the size of each wave has to be calculated using maximum agents and number of influxes. The problem is that due to divisions resulting in fractions (and agents must be whole integers), the accuracy of the estimated migrants to place goes down the higher the number of fluxes is.

Algorithm 2.2 Choose locations for migrant agents during migration waves

```
function PlaceMigrants(p_T, Δ_B)          ▷ Place Δ_B migrant agents around location p_T
    P_T ← {p_T}                           ▷ Initialise the set of locations for immigration
    while |P_T| < Δ_B do                  ▷ More migrants to place?
        p_T ← BestNeighbour(p_T, P_T)     ▷ Find the best neighbouring location
        P_T ← P_T ∪ p_T                   ▷ add it to the result set
    end while
    return P_T
end function
```

Algorithm 2.3 Find the neighbouring site with the greatest population density

```
function BestNeighbour(p_T, P_T)   ▷ Best neighbour of p_T excluding locations P_T
    if |N(p_T) − P_T − {a_i : p_i ∈ N(p_T)}| > 0 then   ▷ Vacant sites not already chosen?
        P∗ ← {}                                  ▷ Initialise best locations
        d∗ ← −∞                                  ▷ Initialise best density
        for all p ∈ N(p_T) − P_T − {p_i : a_i ∈ A_t} do   ▷ All vacant unchosen neighbours
            d ← |{a_i : p_i ∈ N(p)}|/|N(p)|      ▷ Calculate local population density
            if d >= d∗ then
                d∗ ← d
                P∗ ← P∗ ∪ (p, d∗)
            end if
        end for
        return ChooseOneAtRandom({p : (p, d) ∈ P∗ ∧ d = d∗})
    else
        p ←ChooseOneAtRandom(N(p_T))
        return BestNeighbour(p, P_T)
    end if
end function
```

Depending on the rate and size of migration the differences lie between 0.04% and 3.48% of density. The threshold that should not be exceeded is 0.24%, which is less than a quarter of density different and can be rounded down. I.e. 87.24% density is rounded down to 87%. If the actual density differs by more than that, the missing agents are added. To avoid the case in which one influx is larger than the others, all remaining migrants are randomly distributed across all influx waves. In the case of $E = 100$, this can mean that 13 or 14 migrants will arrive at once, rather than 12.

Arrival locations: FluxType

The second round of treatments varies the mode of arrival for migrants, the $FluxType$. There are three kinds of $FluxType$, named $Scatter$, $Cluster$, and $Diaspora$ for readability purposes. $Scatter$ places migrants on random empty locations throughout the map. $Cluster$ will ensure a clustering of migrants around dense areas, including natives. $Diaspora$ will cluster migrants around migrant-dense areas, excluding natives.

I recall that the influx mechanism described in the previous section is the $Cluster$ approach, which seeks out high-density areas regardless of whether these areas are populated by migrants or natives. This reflects our knowledge that migrants are likely to target job-rich environments, which are more likely to occur in more densely populated areas such as cities. Also, there may not be existing diasporas when migrants arrive; diasporas might be too small and thirdly, not all migrants will automatically seek out 'their kind' (Collier, 2013).

Scattering migrants will seek random empty locations on the grid and migrate to these places. Intuitively, the scattering should reduce the impact of migration waves on the native population, because neighbourhoods don't find themselves suddenly outnumbered. Most natives will only encounter a few new migrant neighbours that way.

The diaspora migrants will only cluster around existing migrants. If no migrants are present, the incoming migrants will cluster around dense areas instead. This $FluxType$ is based on existing real migration patterns: some groups of migrants are highly diaspora-oriented and will only be attracted to migrate in the first place if a diaspora exists in the future host country (Collier, 2013), see also (Novotny and Hasman, 2015).

Both *Cluster* and *Diaspora* treatments will treat the incoming migrants *during and after migration* as appealing, but only the *Cluster* treatment will consider natives as appealing during the placement of incoming migrants. There is a distinction between incoming migrants and existing migrants. Because migrants are placed sequentially, the second migrant will know where the first migrant is going to settle. Because this sequential placing happens between ticks, the placement is still 'instantaneous' in that agents were unable to act while migration occurs. Table 2.6 provides an overview of the different priorities of the influx mechanisms.

Table 2.6: The *FluxType* treatments and which circumstances are deemed appealing when placing migrants

Treatment	Incoming migrants appealing?	Existing migrants appealing?	Existing natives appealing?
Scatter	No	No	No
Cluster	Yes	Yes	Yes
Diaspora	Yes	Yes	No

In order to ensure that migrants will arrive in some form of a cluster, both *Cluster* and *Diaspora* will consider the incoming migrants appealing. The *Scatter* treatment does not do this.

Figure 2.5 visualises each kind of *FluxType*. As in Figure 2.4, cyan-coloured tiles denote 'flagged' tiles for future migrant placement. The three different cyan-coloured patterns demonstrate the difference between each *FluxType*: scattered on the left (2.5a), drawn towards dense areas in the centre (2.5b), and the large bloc that 'swallows' existing native areas on the right (2.5c). The bottom row displays the end of the tick at which the migrant wave has occurred: migrants are now placed, and all agents have acted. The scattered migrants show no visible clusters yet. The clustering migrants have caused the existing natives to form medium-sized clusters. The diaspora migrants remain in a moving and shifting bloc, whereas natives on the same map only form smaller blocs.

Figure 2.5: Screen captures of the three kinds of migrant arrival. The top shows the marked places for arrival; the bottom screen shows migrants arrived and having moved.

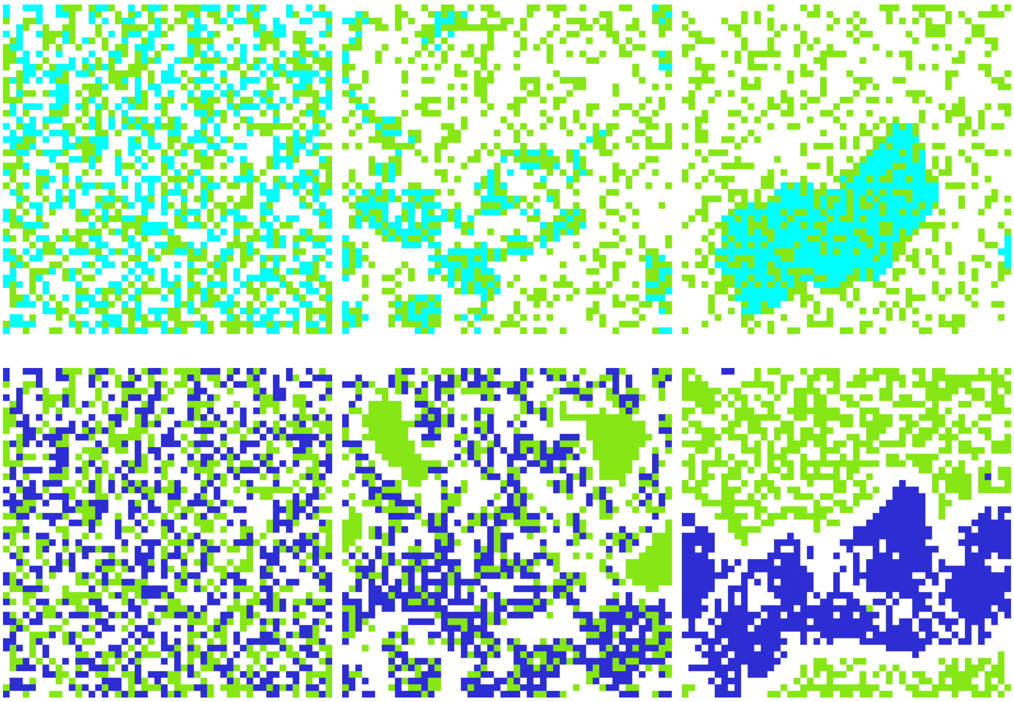

(a) Scattering of migrants at the time of arrival (top) and after moving once (bottom). White tiles are empty space.

(b) Clustering of migrants around densely populated (native) areas at the time of arrival and after moving once.

(c) Clustering of migrants only around other migrants (disapora) at the time of arrival and after moving once.

The *scatter* mechanism will select random empty tiles from across the grid and place migrants there. The *diaspora* mechanism uses algorithm 2.2 for placing migrants just as the cluster *FluxType* does, but it employs a variant of the *BestNeighbour* algorithm 2.3 which prioritises migrants, as shown below, *BestMigrantNeighbour* 2.4:

The model terminates at $t = 2,000$ ticks, meaning that migration takes place in 30-40% of the time, between $t = 200$ and $t = 1000$, depending on the number of migration waves. The time scale of a Schelling model is usually not defined, but intuitively in this case, the total number of ticks could refer to year of high

Algorithm 2.4 Find the neighbouring site with the greatest migration density

function BESTMIGRANTNEIGHBOUR(p_T, P_T) ▷ Best neighbour of p_T excluding locations P_T
  if $|N(p_T) - P_T - \{a_i : p_i \in N(p_T)\}| > 0$ **then** ▷ Vacant sites not already chosen?
    $P* \leftarrow \{\}$ ▷ Initialise potential locations
    $d_B* \leftarrow -\infty$ ▷ Initialise best density of migrants
    for all $p \in N(p_T) - P_T - \{p_i : a_i \in A_t\}$ **do** ▷ All vacant unchosen neighbours
      $d_B \leftarrow |\{a_i : p_i \in N(p) : c_i = B\}|/|N(p)|$ ▷ Local migrant population density
      if $d_B > d_B*$ **then**
        $d_B* \leftarrow d_B$
        $P* \leftarrow P* \cup (p, d_B*)$
      end if
    end for
    return CHOOSEONEATRANDOM($\{p : (p, d_B) \in P* \wedge d_B = d_B*\}$) ▷ of all potential location, pick a random one
  else
    $p \leftarrow$BESTNEIGHBOUR(p, P_T) ▷ If no migrant areas are found, revert to density
    return BESTMIGRANTNEIGHBOUR(p, P_T)
  end if
end function

migration, or a decade of moderate migration, or several decades of little migration. The ratio of migration to the overall time is more important, as it determines the reaction time of agents and frequency of change.

2.2.4 Dependent variables

For each realisation of the model I sample and record dependent variables every 10 time steps, allowing for both cross-sectional and time-series analysis. These variables are described in turn below, and summarised in Table 2.4 at the end of this section. The dependent variables are:

1. **Global happiness** $\bar{h}$: the average happiness of agents

    1.1 **Native happiness** $\bar{h}^G$: the average happiness of native agents

    1.2 **Migrant happiness** $\bar{h}^B$: the average happiness of migrant agents

2. **Segregation of colour** M^c: the Moran's Index of spatial autocorrelation of agent colour c_i

3. **Segregation of tolerance** M^F: which computes the Moran's I not of agent colours but their F value, F_1 and F_2.

The global happiness metric is collected to ease the comparison to many Schelling models that use the agent's happiness as the primary goal of the simulation.[5] It can demonstrate the upheaval that a new influx can cause. I define global average happiness $\bar{h}$ at time t as:

$$\bar{h_t} = \frac{|a_{i,t} \notin D_t|}{|A_t|} \tag{2.5}$$

where D_t is the set of dissatisfied agents (see equation 2.1). To separate the effects of migration on the happiness of both population groups, I also record the average happiness of each group; for natives:

$$\bar{h_t^G} = \frac{|a_{i,t} \notin D_t : c_i = G|}{|a_{i,t} : c_i = G|} \tag{2.6}$$

[5]In several Schelling variations, when all agents are happy, the simulation terminates. This is not the case for this model. The simulation will terminate when t_{max} is reached.

and for migrants:

$$\bar{h_t^B} = \frac{|a_{i,t} \notin D_t : c_i = B|}{|a_{i,t} : c_i = B|} \tag{2.7}$$

As implemented by Benenson and Hatna (2011), I record Moran's index of spatial autocorrelation in order to quantify the amount of segregation by F_1. The Index is a measure of how clustered the agents are. High levels of Moran's I indicate high levels of segregation.

$$M_t^c = \frac{|A_t|}{\sum_{(i,j)\in A_t^2} w_{i,j}} \frac{\sum_{(i,j)\in A_t^2} w_{i,j}(c_i - \bar{c}_t)(c_j - \bar{c}_t)}{\sum_{i\in A_t}(c_i - \bar{c}_t)^2} \tag{2.8}$$

where the mean colour is $\bar{c}_t = \sum_{i\in A_t} c_i/|A_t|$, and $w_{i,j} = 1$ if and only if agents a_i and a_j are immediately adjacent on the lattice (including diagonals), otherwise $w_{i,j} = 0$. I also compute the Moran's I of *tolerance*, M_t^F, where the mean tolerance is $\bar{F} = \sum_{i\in A_t} F_{1,i}/|A_t|$:

$$M_t^F = \frac{|A_t|}{\sum_{(i,j)\in A_t^2} w_{i,j}} \frac{\sum_{(i,j)\in A_t^2} w_{i,j}(F_i - \bar{F}_t)(F_j - \bar{F}_t)}{\sum_{i\in A_t}(F_i - \bar{F}_t)^2} \tag{2.9}$$

A chequered chessboard would have a segregation index of 0, as every white tile is followed by a black tile. A completely segregated grid has an index of 1. Schelling models have a baseline segregation level that is caused by random movement of agents: at any point in time, even a completely random agent with no preferences can, by chance, reside next to a like-coloured agent. This noise-induced segregation level is usually above 0.1 and below 0.2.[6]

[6] The value was obtained by running N=1000 simulations of the parameter range with no agent rules other than random movement throughout the grid.

2.3 Results

This section will describe the model outcomes and explain why such outcomes occur. The implication of the results will be discussed in the discussion Section 2.4. For each migration treatment of $E > 0$ and $FluxType$, the parameters are randomised. Treatments were repeated for $(4E \times 3FluxType) = 12 \times 1,500 = 18,000$ simulations. The control treatment $E = 0$ does not have a $FluxType$ as no migration takes place. The control was repeated $\sim 4,000$ times, resulting in a total of $22,000$ simulations. For every simulation, I record the results at $t_{max} = 2,000$ for cross-sectional analysis, and sample every 10 ticks for time-series data.

I will begin with an overview of model outcomes when $NatShare = 50$ and $FinalDen = 98$, in line with the previously mentioned Schelling models. I will then move on to analysis of varying these parameters.

Figure 2.6 shows boxplots of the dependent variables by each value of F_2 across all E and $FluxType$. I recall that F_1 values are always low and thus agents are tolerant. Both natives and migrants react with lower happiness levels when the F_2 agents are less tolerant. Most F_2 values are above the critical value $F_2 = 17$ but they drive happiness down as F_2 increases. Natives are more negatively affected compared to migrants when $F_2 \geq 21$ and at $F_2 = 20$, mean happiness is slightly higher than when $F_2 = 19$. For migrants, the decline in happiness as intolerance values increase is more gradual. Both population groups have near-total happiness at the critical value of $F_2 = 17$.

The segregation of colour M^c shows that when F_2 increases, segregation decreases. It is significantly higher at the cut-off value of $F_2 = 17$. Combined with the happiness pattern, this is indicative of agents being *too* intolerant to find acceptable neighbourhoods. As they remain unhappy and move frequently, they break up segregation patterns. Agents on the fringes of clusters will always try and relocate as

Figure 2.6: Boxplots of the dependent variable by the F value of the intolerant group at t_{max}. Bold red bars denote means, cyan bars denote medians.

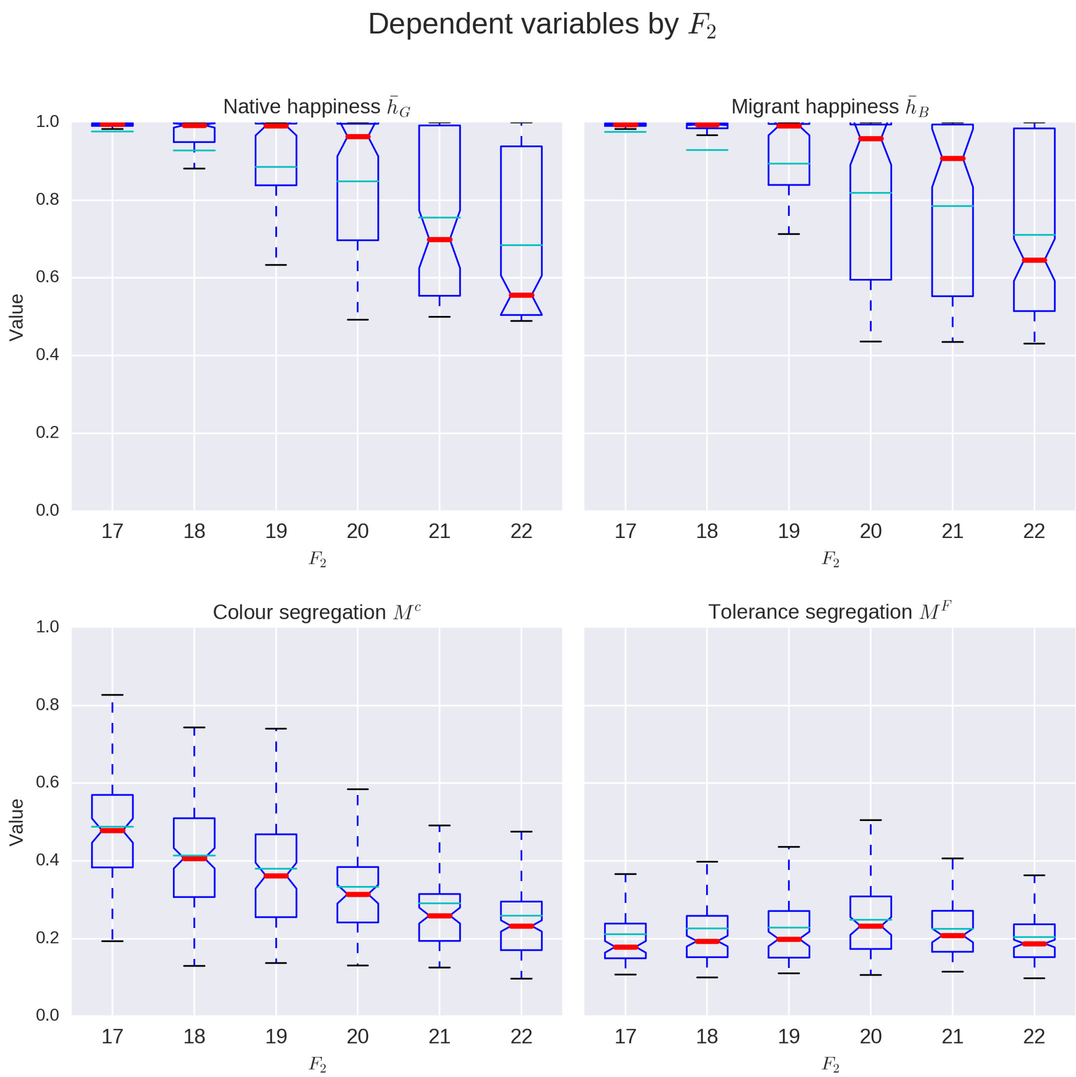

they are in contact with out-group members. The segregation of tolerance increases slightly when $F_2 > 17 < 21$, but overall levels are low. M^F does not correlate with M^c, suggesting that the agent behaviour that drives M^c does not affect M^F in the same way.

To understand why we see such happiness and segregation patterns, I recall the close link between segregation and agents' utility-satisficing behaviour. Segregation is the mechanism with which intolerant agents achieve happiness. We see lower segregation values at $F_2 \geq 20$ because the intolerant agents are *too* intolerant to find satisfactory locations for sustained periods of time. This results in an increase of movement, as perpetually unhappy agents keep moving to achieve the desired state of happiness. In the process of moving away, they lower the ratio of friends for other intolerant agents in the area, creating a domino-effect. Other intolerant agents that were happy become unhappy, and move away. The volatility of this constant movement makes it difficult for neighbourhoods to settle.

Happiness of agents is therefore dependent on moderate segregation levels- provided that the intolerant group has an F_2 value of ≤ 17. This means that for a population to appease its intolerant half, this intolerant half has to be largely segregated. Important to note here is that up until $F_2 = 19$, the vast majority of agents is happy, despite an average segregation value of $M^c < 0.4$. So while some intolerant agents are unhappy, others have achieved a state of happiness despite their high demands on the neighbourhood composition. Because F_1 and F_2 groups are half of the population, an average happiness level of $\bar{h_G} = 0.5$ would indicate that the entire intolerant half of the population is unhappy. Any values above suggest that intolerant agents manage to find happiness.

Figure 2.7 graphs the happiness of all agents over time for the first $1,000$ ticks, comparing the different treatments of E to the control at $NatShare$ at 50%. With no migration, happiness reaches convergent levels at $t = 250$. The impact that each migration wave has is clearly visible as happiness drops at each new migrant wave arrives. The one-off treatment $E = 1$ results in large shocks to happiness, dropping from .95 to .53 within the one turn in which the migration wave arrives. As quickly as happiness falls, it recovers almost instantaneously for $E = 4$, $E = 15$ and $E = 100$ so long as the final waves have not arrived. Past $t = 700$, when $E = 4$, the fourth and final wave of migrants arrive- this time, happiness does not recover as quickly, and it does not reach the same level again. For $E = 15$, this effect starts with the penultimate wave at $t = 789$.

Figure 2.7: The global happiness over time, comparing all migration treatments and the control over the first 1000 ticks when $F_2 = 17$.

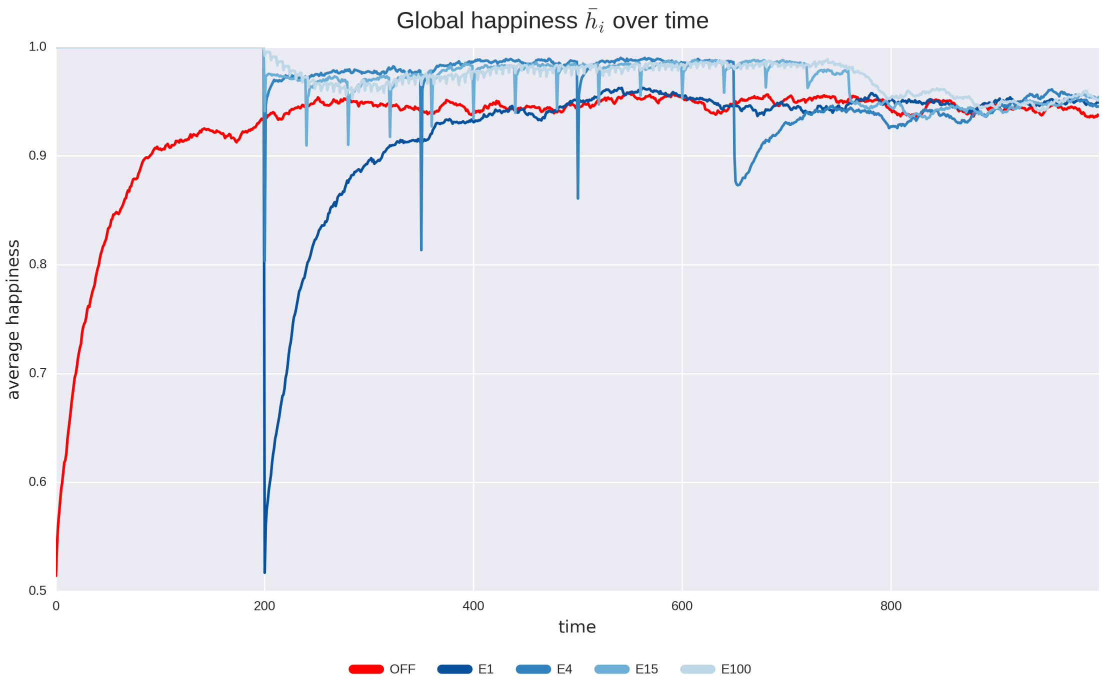

However, when we examine migrant and native agents separately, we see very different patterns of happiness over time: Figure 2.8 displays the happiness of natives overtime, and Figure 2.9 the happiness of migrants over the same time period. On first glance, the native happiness resembles the overall happiness levels very closely. All of the influx-driven dips in happiness appear, and convergence towards 800 ticks and later happens as well. The only difference to overall happiness, as shown in Figure 2.7, is that natives in the case of $E = 100$ and $E = 15$ have reached near-total happiness. Every shock to happiness through migration is absorbed within a handful of rounds. The $E = 4$ treatment shows similar trends, despite its much larger downward spikes.

Migrant happiness (Figure 2.9) is much more erratic with the exception of the control $E = 0$, in which migrants behave just as natives do. The happiness levels over time for all other treatments do tend upwards overall and converge at the same levels as native happiness does, but in the periods where migration still occurs, migrant happiness shows a different pattern. At first, each migration wave causes

Figure 2.8: Native happiness $\bar{h}^G$ over time, comparing all migration treatments and the control over the first 1000 ticks when $F_2 = 17$.

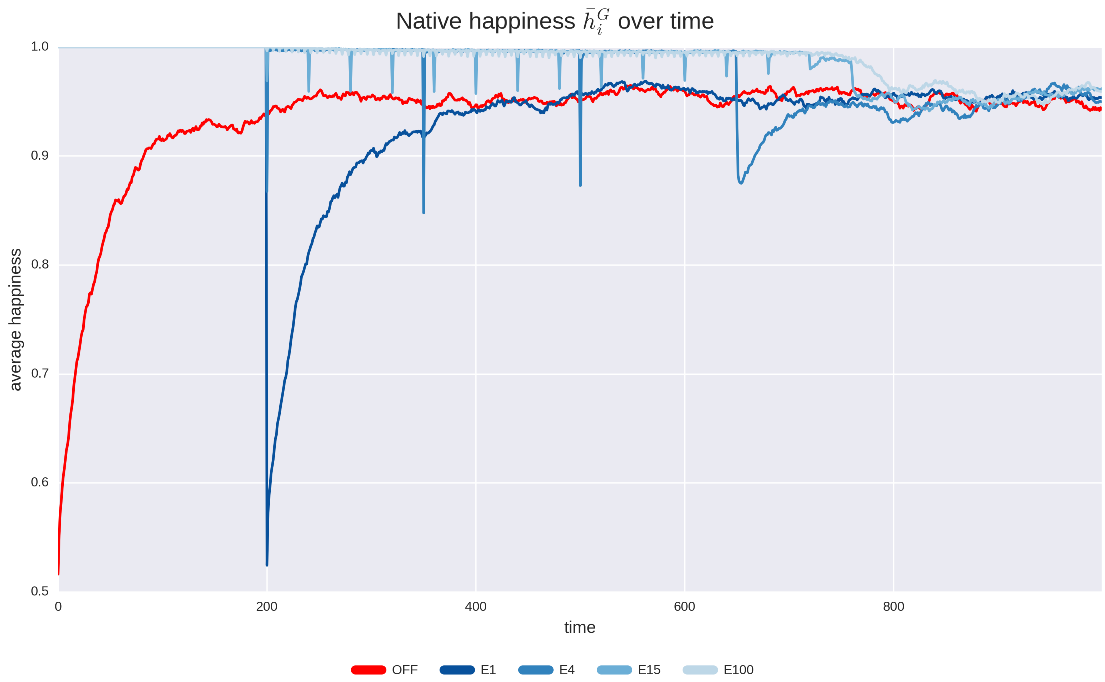

an increase in migrant happiness, before decreasing at various degrees, depending on the treatment. This is clearly visible for the cases $E = 1, 4, 15$, less so when $E = 100$. For influx treatments less than $E = 100$, the next migration wave causes a drop, rather than an uptick in happiness. For $E = 100$, this trend is delayed: each wave still causes upticks up until $t = 400$, after which subsequent wave cause a down-turn in happiness of migrants, too.

The changes in happiness are largely driven by two things: the proportion of the population that each group represents, and the fact that happiness of new agents (be it native or migrant) is randomly initialised. The population proportion explains why the overall happiness levels are so sharply dominated by natives. For many treatments, they constitute a large chunk of the total agent population, and the overall happiness levels do not capture the differences within each group. The random initialisation of happiness is visible when looking at the $E = 0$, or OFF treatment, on any of the graphs. Because agents are initialised at the start, the average happiness in the first round is always around 0.5. Combined with the proportion

Figure 2.9: Migrant happiness $\bar{h}^B$ over time, comparing all migration treatments and the control over the first 1000 ticks when $F_2 = 17$.

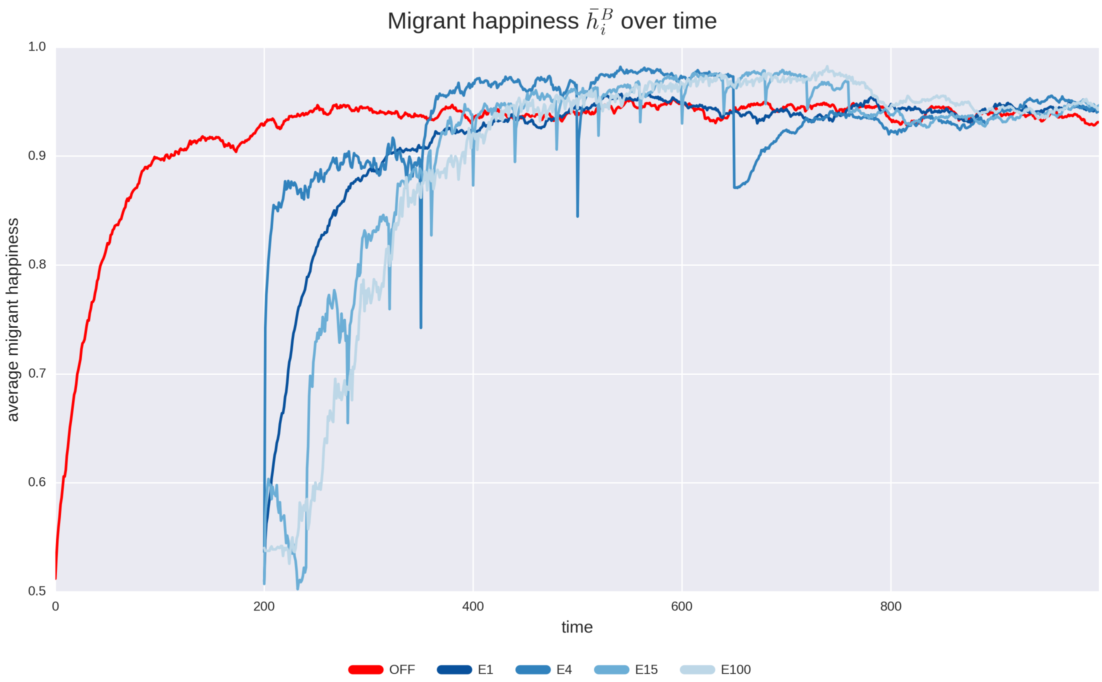

size, this explains why migration waves first seem to increase migrant happiness, and then decrease it. When migrants are few and happiness is still below the average at the time, each new wave comes with roughly half happy, and half unhappy migrants. These push the overall happiness up. As average migrant happiness increases over time, the unhappy new migrants drive down the overall levels. The drops in migrant happiness are thus 'artificial' in the sense that a proportion of migrants is always pre-determined to be initially unhappy. Existing migrants cannot experience unhappiness as a result of new migrants, as they consider them friends.

Figure 2.10 below plots the segregation over time at the critical value of $F_2 = 17$. As the boxplot above indicated, long-term convergence of the Moran's I is around .48 for all treatments and the control. The migration waves cause shocks upwards as segregation increases. When $E = 15$, the first wave causes a 1-tick upward shock, but then declines steadily until the next wave arrives and sends segregation upwards once again, After the arrival of the third wave, the falling segregation levels are less pronounced, and the effect vanishes with further migration waves. When $E = 100$, the same effect can be seen at the first couple of waves at $t = 200$, but the differences are subtle. This suggests that when $E \geq 15$, enough migrants arrive at once to form a cluster, but when migration waves are too small, migrants cannot create a stable cluster. I recall that always 50% of each group is F_1, so when only 20 migrants exist of which 10 are happy living around natives, the quantity of familiarity-seeking migrants is too low.

Figure 2.10: Colour segregation over time, comparing all migration treatments and the control over the first 1000 ticks. Higher values indicate higher levels of segregation of green and blue agents.

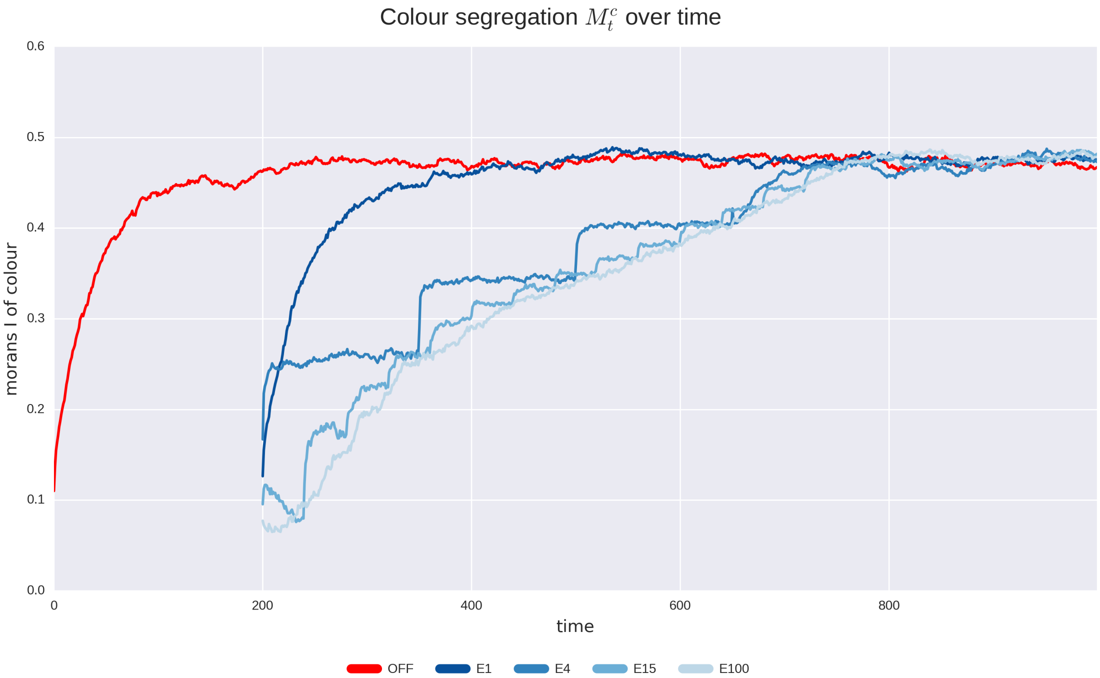

The control treatment has consistently higher values of both segregation and happiness, but towards the end, the values of both $E = 0$ and $E = 1$ converge to very similar levels. The $E = 4$, $E = 15$ and $E = 100$ treatments show the same pattern: every time an influx occurs, the system experiences an overall shock relative to the size of the influx. The small and large step-like increases are very visible in the segregation patterns. As with the $E = 1$, all values eventually converge on very similar and often-times overlapping levels, showing no distinct difference over the long term. In other words, how often migration happens and how big each wave is, does not change the circumstances in the long run. As density approaches 98%, the model outcomes are virtually indistinguishable.

In the short term however, differences are big; in particular, the $E = 1$ experiment stands out. In that setting, the target density is reached in an instant, and thus the behaviour converges to similar levels of the $E = 0$ density of 98%. The convergence takes nearly 200 ticks to occur (between 200 and 400), which means that the short-term impact of a one-off migration wave causes lower happiness levels than when migrants existed from the start. The results indicate that the size and rate of migration in this model does not alter segregating behaviour or outcomes. The Schelling-esque model patterns emerge as soon as critical densities are reached.

Because the convergence of behaviour occurred in the latter stages of the experiments, I zoomed in on the latter stages of the first 1000 ticks. Figure 2.11 shows a close-up view of the global happiness under Treatment $E = 100$ with small waves of 12-14 migrants each. The happiness is plotted against the board density at the time.

Figure 2.11: Time-series of average agent happiness as the number of agents increase.

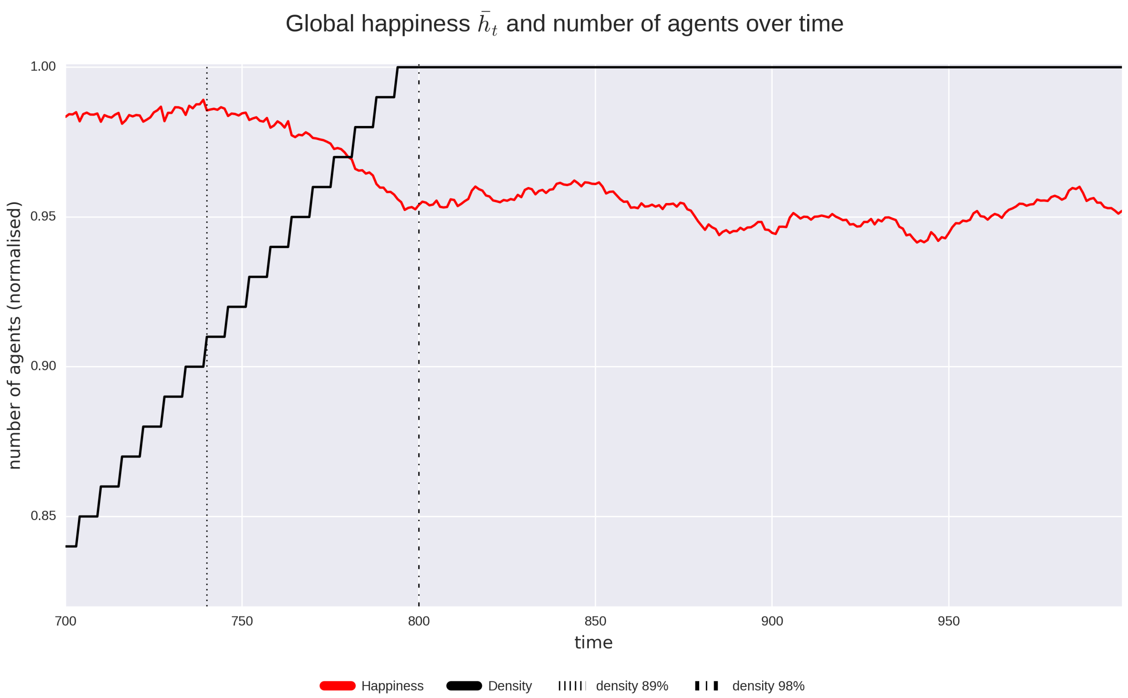

The two vertical lines denote the start and end-point of the drop in happiness until convergence is reached, 89% and 98% density respectively. When the target density is lower than 89%, the convergence of behaviour would not occur and the happiness level would have remained 3% higher. This confirms the importance of density in a Schelling model, which is robust to migration.

I now turn to the *NatShare* parameter which had been 50% for the previously discussed time-series graphs, and the impact of *FluxType*. Figure 2.12 plots low levels of *NatShare* on the left (a,b,c,d) and high levels on the right (e,f,g,h). Each row represents a *FluxType* and the top row (a,e) is the $E = 0$ control without a specified *FluxType*. The happiness of each population group is plotted in their respective colours, native green and migrant blue.

The population group enjoying their majority is near-total happiness for all treatments. (b) and (c) show the slightly lowered happiness just after the migration wave has entered. As a general rule, natives, when in the minority (b-d) start off happy, and stay happy when the first migrants arrive, and then decline until resting at values between .91 and .96 for the different *FluxType* treatments. When migrants are clustered (b), natives are the happiest in the long run; even diaspora arrival (d) records slightly higher long-term happiness levels for natives compared to scattering migrants (b).

The right-hand side of the plot shows a general pattern of similar behaviour: migrants in the minority are not as happy as the majority-natives. Their happiness starts off very low as they begin to arrive, and moves upwards as more waves enter. Migrants are usually less happy than their native counterparts in a minority, suggesting that the late arrival impacts the long-term happiness. When migrants arrive at diasporas (h), they are consistently happier than their clustering cousins (g) and happier than the scattered migrants (f), especially at the early stages of arrival. At $t = 250$, diaspora-seeking migrants are just below 80% happiness, whereas the scattering migrants and the same time period are at 40% average happiness. This

Figure 2.12: Native and migrant happiness over the duration of 2000 ticks, plotted by $NatShare$ and $FluxType$. The vertical dashed lines indicate the end of migration waves.

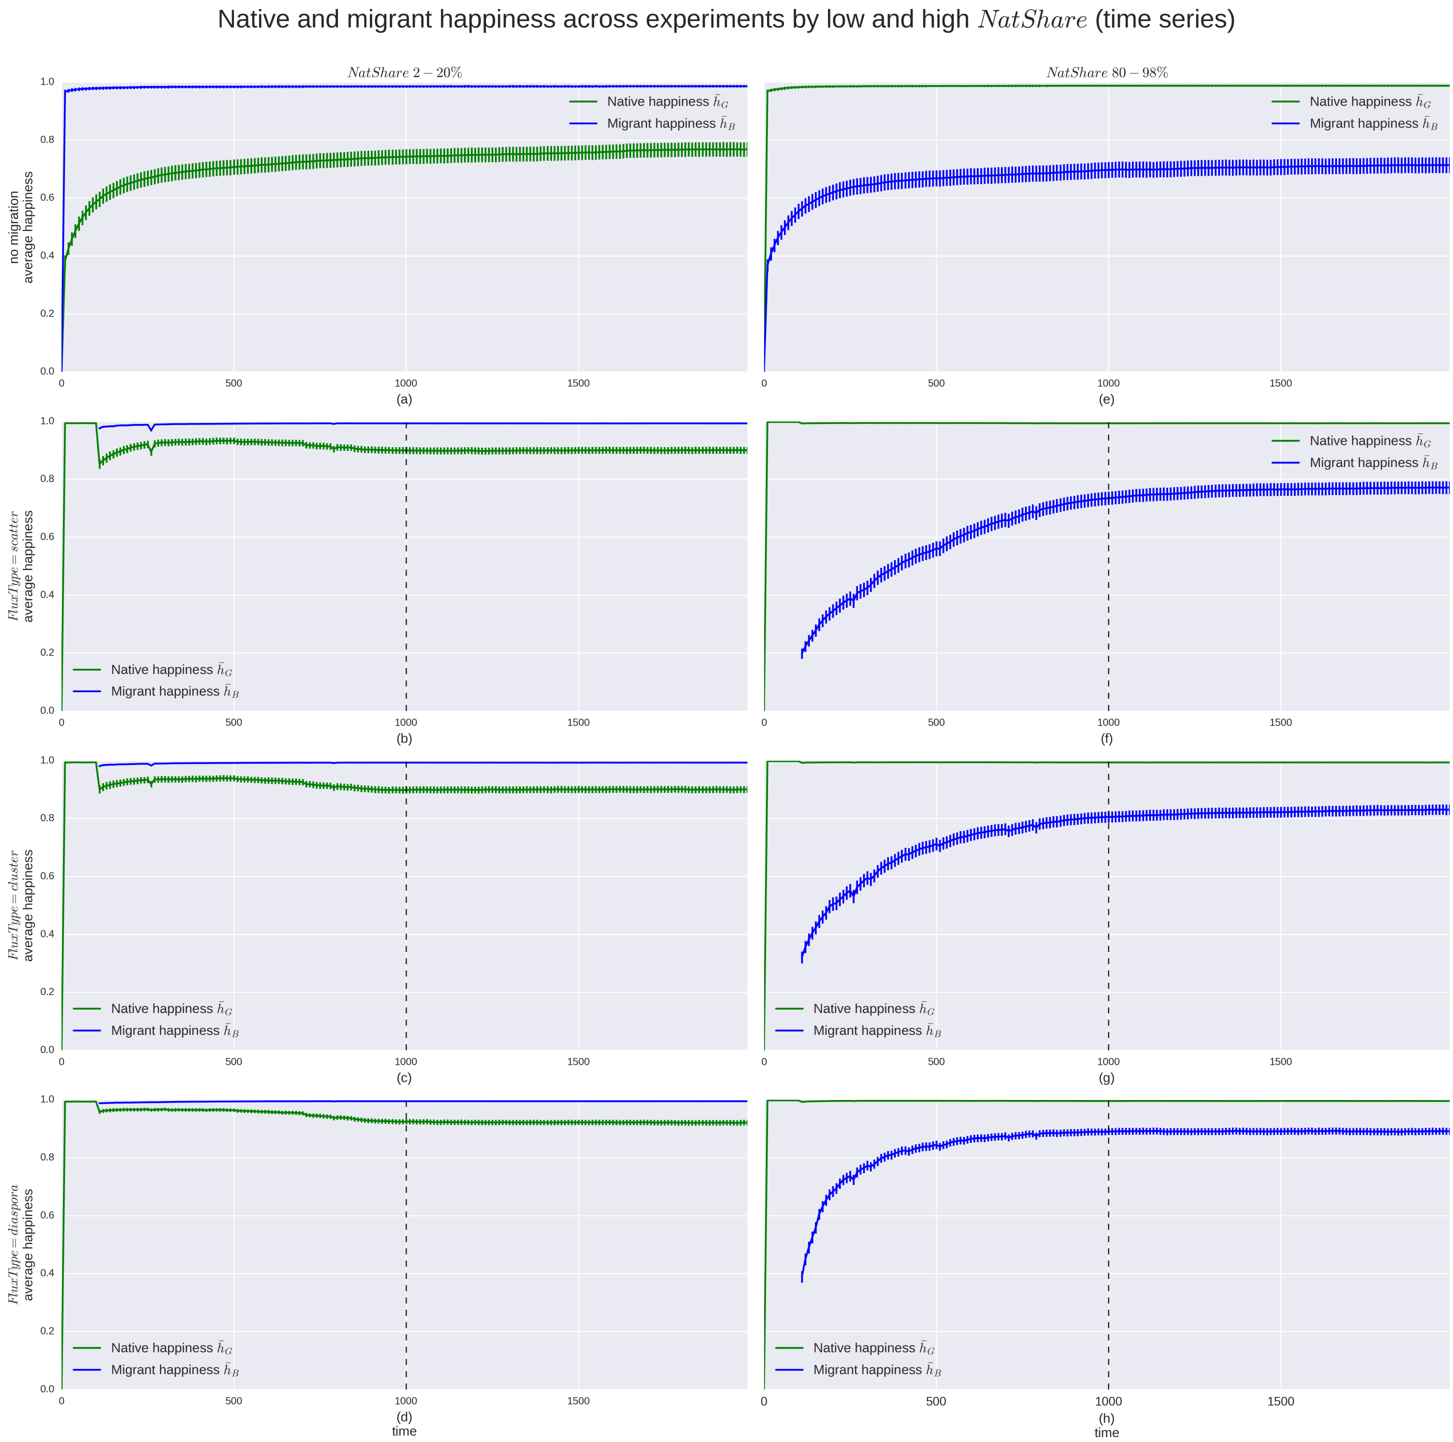

shows that by targeting existing migrant populations, new migrants can achieve higher satisfaction levels especially when they are among the first to arrive. At later waves, the effect diminishes. Clustering migrants (g) are roughly in between the other $FluxType$ treatments on levels of average happiness.

I confirm the differences in behaviour of natives and migrants when each are in a minority of 20% or less in Figure 2.13, which shows the happiness levels at the end of the simulation at t_{max}.

Figure 2.13: Boxplots of native and migrant happiness when each group is in a minority of 20% or less. Bold red lines denote mean, cyan lines denote median values.

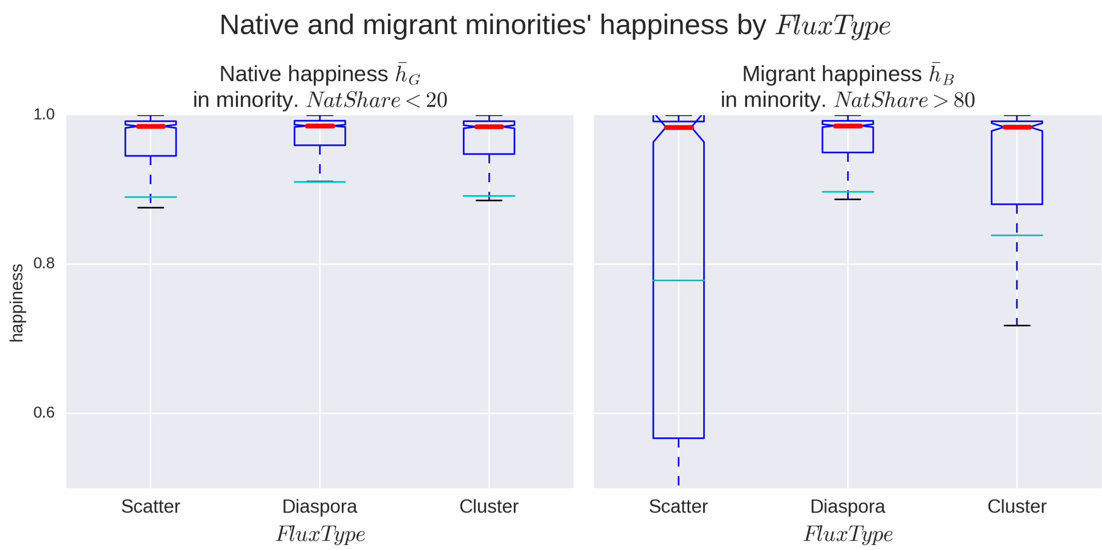

The different $FluxTypes$ impact natives very little, but the subtle increase in happiness under diaspora settlement is visible with an increased median. The boxplot 2.13 shows more clearly than the time-series graph that the scatter treatment has a much higher variation in outcomes as all $E > 0$ treatments are combined. The mean is the same for all treatments, but the medians vary significantly, and the superiority of diaspora-guided settlement as measured by happiness outcomes is visible. The reason that migrants are more strongly affected by changes in $FluxType$ is straightforward: the rule is designed to affect their initial settlement specifically. Natives are only affected indirectly. When natives are in the minority, they show only subtle reactions to the different *FluxTypes*. The Diaspora settlement leads to a slightly increase median happiness. This is because when natives are in the minority, the settlement of the then-majority of migrants inadvertently affects most of natives. The scattering of 80% or more of the population has a high chance of affecting most natives on the grid. Equally, the clustering of such a large proportion of the population around a high-density minority neighbourhood can overwhelm the local population. The diaspora settlement avoids this issue by purposefully settling the migrants away from the natives, resulting in the slightly higher median of native happiness. But as previously noted, these changes at this stage of the model are minor.

Major changes result in the happiness of migrants when they are in the minority. Most notably, the scatter settlement results in a large variance in Migrant happiness $\bar{h_B}$, despite the mean being close to 1. Note the median which is fairly low, just below $\bar{h_B} = 0.8$. This suggests that whilst half of the migrants have happiness levels below that value, the vast majority of happy migrants are *very* happy, thus keeping the mean at near-maximum. Migrants that are scattered have thus very little way of dealing with being dealt a bad set of cards: if a scattered migrant ends up in a largely-native neighbourhood, and the migrant is intolerant, the migrant is unhappy and has to start moving to find the right place to live. Because the tolerant migrants who have ended up, by chance, in a majority-native neighbourhood are happy, they don't move away, thus lowering the number of potential neighbourhoods for intolerant migrants to cluster into. Scattering migrants hurts the intolerant migrants, as the tolerant ones don't provide a buffer in which intolerant migrants can settle. The clustering settlement shows this effect in diminished strength: the vast majority is very happy, but the number of not-that-happy migrants has decreased. Because migrants cluster, only the fringes of the cluster will be effected: this can occur when the local population density is too high to establish a neighbourhood connection between two migrants. The next vacant tile might not be in visible range of another migrant, despite it being the closest vacant tile to the cluster. Intolerant migrants find this situation unacceptable. However, the clustering mechanism greatly reduces the number of unhappy migrants compared to the scattering $FluxType$. The diaspora $FluxType$ achieves even better happiness levels because the problem of high local densities is circumvented by picking less populated spots to cluster around.

$FluxType$ can change happiness, but has no impact on segregation outcomes, as shown in Figure 2.14: again, low $NatShare$ is on the left, high $NatShare$ on the right, and each $FluxType$ on a row with $E = 0$ at the top. The tolerance segregation M^F is convergent and no different from the control. The colour segregation M^c is slightly higher in (b-d) compared to its control: migration increases segregation but the arrival type does not impact this pattern any further. The same trend is visible for high $NatShare$ levels. Graphs (f-h) show slightly elated values of M^c at .6 compared to .55 under no migration (e). I can conclude that colour-segregation behaviour is not impacted by $FluxType$, but increases with $NatShare$ and with F_2. Colour segregation M^c_t is lower when $E = 0$. This is because when no migration occurs, intolerant and tolerant agents have an equal likelihood of meeting outgroup members at the beginning of the simulations. This leads to hardliners segregating, and tolerant agents forming buffer zones by remaining on the outskirts of clusters or not being in a cluster in the first place. When $E = 0$, hardliners are exposed from the start. When $E > 0$, the tolerant agents have clustered with intolerants, forming larger clusters containing both tolerant and intolerant agents. When migration waves come in now, intolerant agents on the fringes of clusters are caught off-guard. Because tolerant agents can reside inside intolerant clusters (they don't mind either way), it creates fewer opportunities for intolerant agents to be happy with the current situation. Thus, hardliners will try and form clusters. The movement patterns settle down when clusters have an intolerant core surrounded by tolerant buffer. This mechanism also explains why colour-based segregation is higher when natives are the vast majority of the population. It is them who exist before the first wave of migration and that create clusters without providing support for intolerant natives. Fringes thus become overexposed and start breaking up the previously established clusters.

Figure 2.14: Colour and tolerance segregation over time until t_{max} by E and low and high levels of $NatShare$. The dashed lines represent the end of migration waves.

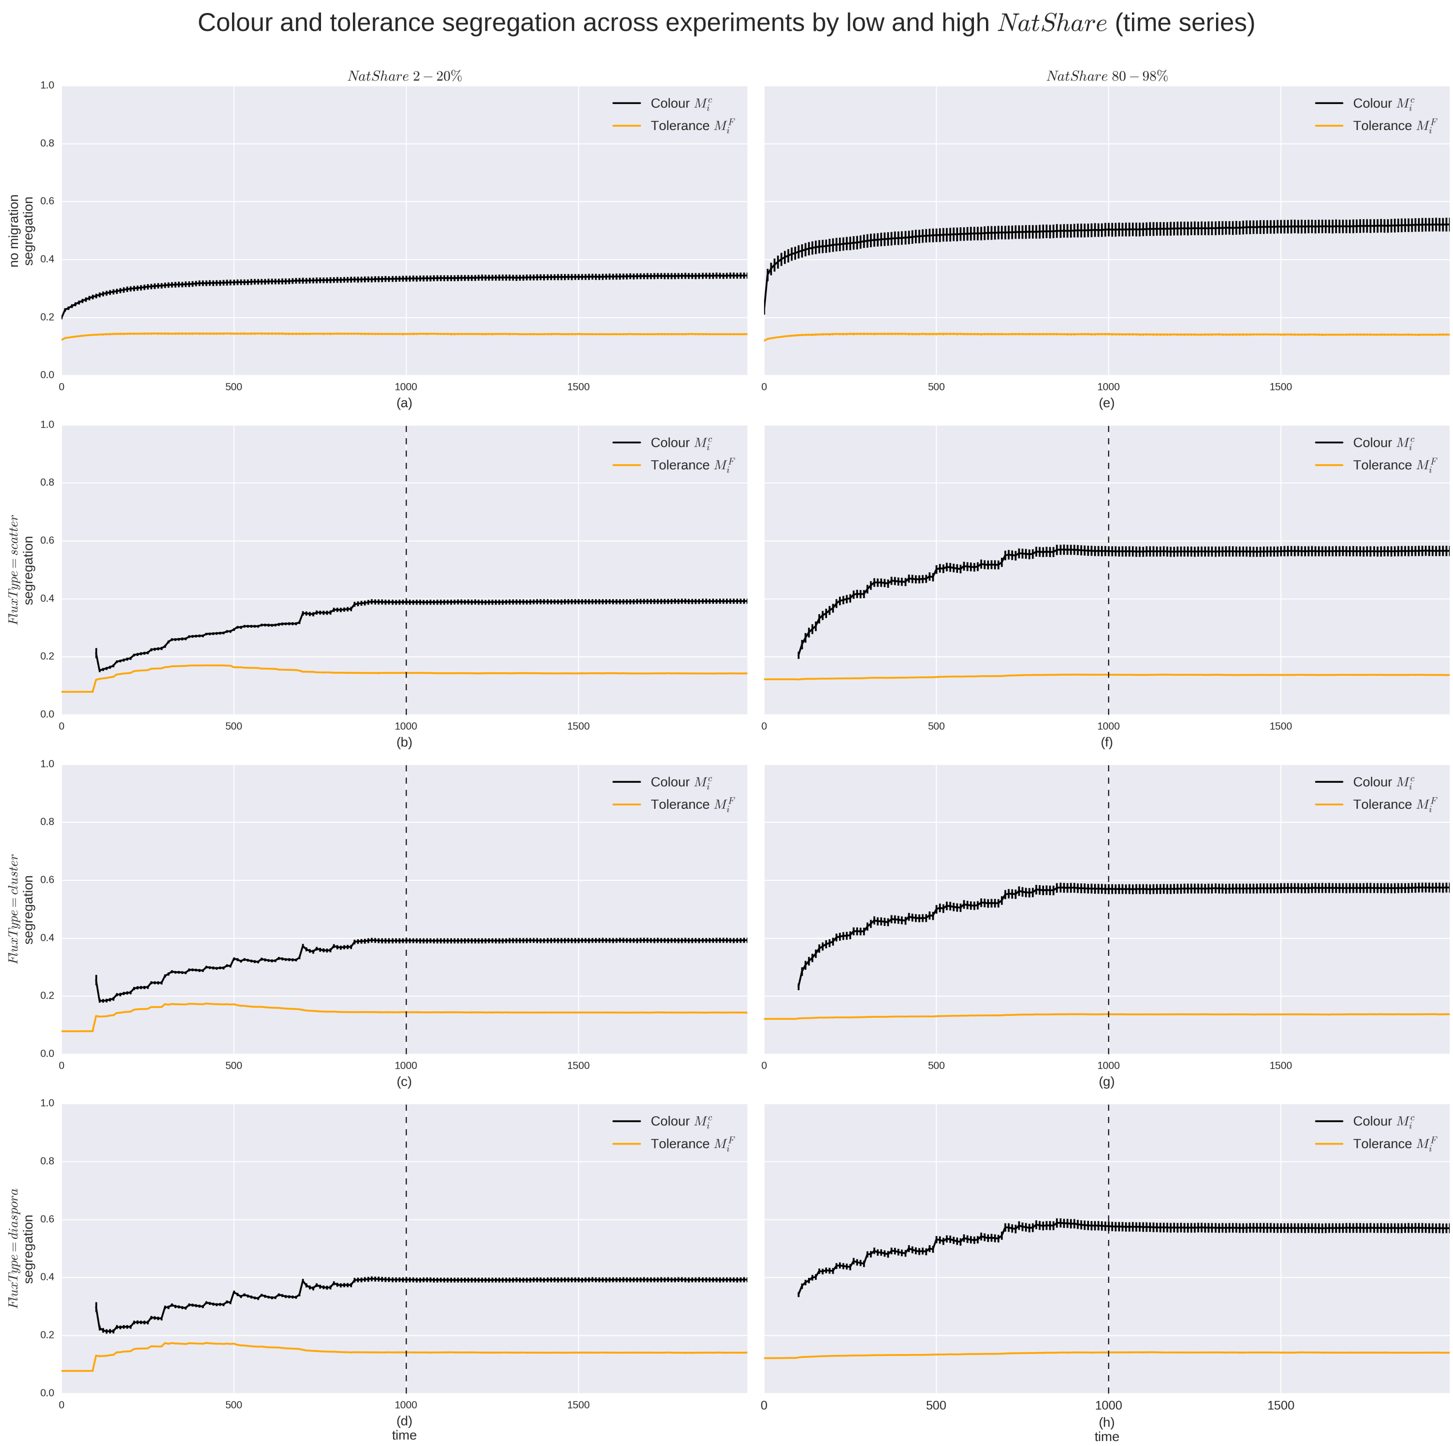

Colour segregation is consistently higher by up to 20% when $NatShare$ is high. When migrants form the large majority, migration is lower overall and over time. This shows that the way agents are introduced does alter long-term outcomes. Not of happiness, but of segregation. When migrants arrive in waves forming a majority which grows larger over time, the society is overall less segregated without sacrificing happiness. Because the natives are so quickly outnumbered, they effectively have to look for the highest possible diversity situation when density is still low and the range of options are high. This allows agents to find satisfactory neighbourhoods. As the migrant majority grows, they may not need to move at all, as they already reside in the right place. In the real world, migration generally involves a context in which the natives-born or previously resident people form a majority. Whilst the majority might shrink, it usually doesn't shrink to percentages as low as 2% or 20%.

2.4 Discussion

The model in this chapter sought to examine the impact of migration on the host society using a Schelling model with very few moving parts. The short-term impact of migration differs significantly from long-term impacts in this model. The contextual neighbourhood effects are dependent on density and the ratio of migrants to natives.

The *size and rate of migration* of a new set of agents into an existing population have no long-term impact on happiness levels of agents. Assuming an equal population share between natives and migrants, any short term variation that may exist between large and small migration flows can be explained by the differences in population density: higher levels of density correspond to lower levels of happiness, because agents have fewer options to move to a satisfactory place. Free space in a Schelling model is a proxy for freedom of choice of residency within towns and cities.

The lack of long-term impact corresponds to the general findings of the existing literature and supports the concept that given passage of time, host societies *absorb* first-generation migrants (Collier, 2013). However, the absence of such ramifications in this model do not prove that such ramifications cannot or would not occur in the real world. It shows rather that mere segregating behaviour alone wouldn't, at least in the absence of mediating factors or behaviours, lead to long-term negative effects. This finding is in line with our knowledge of long-term reactions to migrants. After several decades, new types of migrant groups arrive that shift the framing of difference of existing migrants (Hatton, 2005). Changing discourse and framing can re-cast different population groups, such as post-9/11, which lead to the primacy of Islam as the defining attribute that distinguishes a population group from its (Western, mostly Christian) group (Joppke and Torpey, 2013).

The impact of scale and size of migration in the long run only matters under conditions of lower density, where overcrowding is less likely to cause constant friction. When the number of arriving agents is low, agents have, on average, more empty space to use. Intolerant agents thus seek to belong to groups of their own colour, and if they are on the edge of a group, they are usually surrounded by empty space. Visually, this results in white 'buffer zones' between groups. These zones do not appear if the population covers 98% of the map. Thus, agents on the fringe of a group will likely neighbour agents of a different colour, causing unhappiness and more movement. Because the space to move is so limited, constant friction keeps agents less happy and in a perpetuated state of seeking better places. With more space available, differences in migrant and native population are more visible.

The terminology of 'low' and 'high' density levels however is relative. Whether a figure of 88% density translates into a highly or lowly populated population in the real world will shape any conclusions drawn. Because the Schelling model can resemble urban spaces closely, and urban spaces are confined, high densities in the model can construct the reality in which the housing market restricts the choices of movement (Benenson and Hatna, 2011). Small differences in density can be reminiscent of property development and new opportunities opening up on the housing market: this has been used in various Schelling models of urban geography (Hatna and Benenson, 2012). Previous research in urban development and spatial segregation in cities has shown that the availability of different housing types can drive segregation (Arbaci, 2007). Economic status can also reduce the options and change the preferences of people moving: poor people struggle moving into non-poor neighbourhoods (South et al., 2011).

Aside from the rate of change of migration, different initial settlement strategies were compared. The settlement strategies had a mixed effect on model outcomes. Natives in the model did not react negatively to migrant diasporas, their happiness was even slightly higher in these circumstances. Because the migrating group segregates right away, intolerant agents are effectively shielded from negative experiences. Both groups have their share of tolerant and intolerant agents. When tolerant migrants scatter or target high-density native areas, they may not be compelled to seek out a diaspora because they have settled and are content. This reduces the number of migrants that live in a diaspora, making it harder for intolerant migrants to find satisfactory places to live. When tolerant migrants settle in a diaspora, they are again not compelled to move, and contribute to a larger migrant community that can satisfy its intolerant members. The implications of the impact of migrant settlement types on happiness are that it is in the migrants' interest to seek out their own right away, even if that drives up segregation. In essence, this is what migrants in the real world already do (Collier, 2013) and this model generates the same pattern. The model suggests that happiness-segregation trade-offs are what need to be considered when understanding the impact of migration on hosts and migrants. Initial settlement preferences of migrants may differ from those that develop after migrants have settled in (Søholt and Lynnebakke, 2015).

If diasporas are to be avoided, it is the more tolerant migrants that can be targeted as they will be more open to live in a minority situation. This is reminiscent of high-skilled vs. low-skilled migrant workers who have very different settlement patterns (Joppke and Torpey, 2013). Crucially, migrants that scatter experience 50% less happiness on average than the diaspora migrants. The results suggest that a targeted policy of scattering migrants, designed to reduce individual local impact on natives, might need to consider that the reduced local shock comes at the cost of migrant happiness. In this model this has no further consequences, but in a country that relies on migrant labour, these potential deterrents should be considered. For instance, immigration in Europe in the past decades has branched into two extremes: low skilled (yet useful) labour, and highly skilled migrants that attract niches of the economy (Baganha et al., 2006). Skilled migrants have the resources to chose alternative destinations, and unskilled migrants (particularly those that stay in the country illegally, by means of overstaying visas for example) have ever more reason to "keep to themselves" (Baganha et al., 2006).

Contrary to the long-term convergence, the short-term outcomes in the model are significantly affected by the rate of change of migration. The short-term dip in happiness in a one-off migration situation however is by far the largest and takes long to recover. These short-term differences could be meaningful when considering electoral cycles: even though different migration patterns may theoretically lead to an equally happy society, short-term dips in happiness could be decisive in an election, where long-term prospectives might not be politically useful. This is particularly relevant in recent years which has seen the rise of populism in Europe and the US, partially preying on fears of *Überfremdung* ('over-foreignization') fuelled by the refugee crisis (van Prooijen et al., 2017); but also previous backlashes to migration waves in the 20th century (Hatton, 2005). The refugee crisis has posed (and still poses) a great political challenge for EU countries for parties on all ends of the political spectrum (van Prooijen et al., 2017). The short-term effects are thus substantial and affect future political decisions and ideology. The model suggests that negative responses result through the pure logistics of such a large-scale and short-term change. That is, threat perceptions, identity politics or populist narratives are not necessary to generate shock effects within the population. We know that attitudes are slow to change and that people prefer the status quo for prolonged periods, rather prolonging some discomfort than risk to cause upheaval (Shamir and Shamir, 1997). The sheer size of a logistic challenge should thus be considered, no matter how well-justified it may be on an ideological basis.

These short-term effects matter. Perhaps in fifty years time, the refugee crisis will have been forgotten or overshadowed by different events. But its impact on political forces and on media, society and discourse in this short period time is substantial.

Lastly, the model included different strengths of tolerance to see how different tolerance groups react to migration waves. The grade of intolerance matters. Lower values of high intolerance F_2 were correlated with higher values of happiness in each population group. If the intolerant agents are just a little less intolerant, this can lead to significant differences in overall happiness without driving segregation up as well, as shown in Figure 2.6 when $F_2 > 17 < 21$. Translating this mechanism into public policy might mean that it is not required to 'convert' large swathes of the population to become completely open towards migrants, but that small differences can matter, too. Instead of presenting migration as a pro vs. anti binary political stance, it could be construed as a scale that acknowledges a potentially wide-spread desire to be in the majority population group. The importance of framing has been highlighted in diversity research, especially when reaching out to 'hardliners', i.e. people who self-identify with the majority population and who view their race and ethnicity as central to their identity (Rios and Wynn, 2016). Hardliners are more open to multiculturalism when it is presented as a learning opportunity to learn about other people, rather than as a political aim or ideological pursuit (Rios and Wynn, 2016). Learning about other cultures or ethnic backgrounds can constitute a first contact situation and provides an opportunity to adjust pre-existing prejudices. If this lowered perceived threat (and subsequently, intolerance) just mildly, it could have a positive effect overall.

As a final concluding remark, the long-term convergence speaks to the robustness of the Schelling model. Even introducing 96% of a population at once does not alter the general model outcomes and patterns- intuitively, a one-off shock of a large number of migrants should be more disruptive in the long term. A long-term disruption could be caused by path dependency: because of the initial shock, agents segregate strongly. If these segregation patterns persist, then the initial shock would have altered the long-term outcome. This is not the case with this model.

3 Chapter 3

In this chapter[7], I build on the immigration model introduced in Chapter 2. The aim of this section of the thesis is to implement an existing social theory into an Agent-based model to test the theories' assumptions in a repeated virtual experiment. The theory to be implemented is the so-called 'contact theory' which draws from wider work on group identity in psychology and sociology. By implementing a theory in a computer model, its assumptions and predicted outcomes can be scrutinized. The implementation of the contact theory introduces a crucial aspect into the existing Schelling model of immigration: adaptive tolerance. Depending on the experience with migrants, native agents can now respond by raising or lowering their tolerance towards out-group members. I again address the research question (i): "How does migration affect the host society and migrant population?". The introduction of immigration to a Schelling model tested two very specific concerns relating to migrant settlement and rate of change of migration. The model introduced in this chapter uses existing social theory that makes assumptions about contact between two different groups of people- such as migrants and citizens of the host country.

In the realm of diversity, two main strands of thought have crystallised. The contact theory posits that diversity increases tolerance and solidarity (i.e. social capital). The greater the exposure to otherness in a social environment, the more likely is it that people perceive these differences as the status quo, and they cease to be markers of difference. Opposing the contact theory is the conflict theory (or threat theory). The otherness, so the theory claims, triggers anxieties and leads to a protective reflex. The higher the diversity, the more likely people are to stick to their own kind and try and turtle, mistrusting the other. Both theoretical approaches have empirical evidence to support their claims (Kaufmann & Harris, 2015). Problems arise in both the operationalisation as well as the level of analysis. Capturing

[7] This chapter is based on a paper co-authored by Steve Phelps (KCL): Urselmans and Phelps (2018)

migration on national or sub-national level does not entail data on the individual or group-based interactions, which are the crucial level of analysis for social capital. Kaufmann and Harris (2015) provide a meta-analysis of studies testing the minorities, migrants and outgroup-hostility. Note that the share of minorities and share of migrants are sometimes combined in these studies. Of the studies listed, 11 support conflict theory and 8 support the contact hypothesis. Two studies find a neutral relationship, while the remaining 4 report multiple responses (i.e. both negative and positive, or neutral and negative). The absence of consensus on contact or conflict is evident. The authors point at the difference in level of analysis as a crucial factor in deciding the outcome: movement is easier to capture with smaller units of analysis. Ariely (2014) echoes a similar observation, analysing studies conducted between 2007 and 2014. Especially the operationalization of variables such as trust leads to contradictory results between studies, even if studies are internally valid.

I bridge this micro-macro gap by employing an agent-based model suited to explore theories of social complexity through the ability to capture the properties of heterogeneous populations of individuals (Klabunde and Willekens, 2016b). This bottom-up approach in which the model is imbued with micro-level behaviours, can give rise to macro-level behaviours which can be observed empirically in the output from the model; thus it is easier to bridge the micro-macro gap because all micro-level behaviour is automatically accounted for.

3.1 Introduction

International migration is becoming an increasingly defining feature of Western countries, and a key question is what the social implications of large-scale migration and increasing ethnic diversity are. As societies grow more diverse and immigration increases as the world becomes increasingly globalised, concerns about immigration in Europe are high (Heath et al., 2016). The parsimonious model introduced in Chapter 2 revealed the importance of minority population share and grades of intolerance, but is still based on several unrealistic assumptions about the real world. To bring the model closer to corresponding literature on intergroup contact, this chapter implements the contact theory to test its micro-foundations.

That people harbour (even slight) preferences for belonging to social majorities is well-known; disagreements occur when trying to explain why these preferences exist and how they come into existence (Singh et al., 2009). One such explanation is provided by group identity theory (Clark and Fossett, 2008). People will self-categorize themselves and others into 'us' and 'them' and establish differences between these categories as a means of identifying with a group and to experience the benefits of belonging to a group (Pettigrew et al., 2011).

In social psychology, immigrants and citizens of host societies can be understood as two different groups that lend identity to its members. To native-born people, immigrants thus form an 'outgroup'. Intergroup threat theory describes the perceptions of threat that people perceive from an outgroup (Stephan et al., 2009). Perceived threats to society and culture from immigrants are worrying large amounts of voters across Europe (Heath et al., 2016). Threat theory includes perceptive threats, which will be the focus in this chapter. Thus, whether or not the threat is real is not the primary concern: what matters is that people feel *as if* it were real. Survey respondents frequently overestimate the number of immigrants in their country (Markaki and Longhi, 2012). Actual numbers of migrants do not predict perceived threat (Semyonov et al., 2004; Stephan et al., 2009), but perceived numbers do (Semyonov et al., 2004). This can help explain why anti-immigration attitudes are often high in areas with low migration: following the Brexit referendum in 2016, Goodwin and Heath (2016) have examined the demographics of voters. They find a negative relationship between EU migration and support for leaving the EU: "of the 20 places with the most EU migrants 18 voted to remain. In many of the areas that were among the most receptive to the Leave campaign there were hardly any EU migrants at all." (Goodwin and Heath, 2016, p.10).

Intergroup threat theory has proven an effective tool in testing what drives such sentiments (Stephan and Renfro, 2002). Perceived threats to key values in society can explain the anti-immigrant hostility in Europe (Mclaren et al., 2012). Those who believe that traditional values are undermined and that societal cohesion is not 'what it once was', are also more likely to be sceptical of immigration (Mclaren et al., 2012).

Threat theory has drawn elements from intergroup contact theory, initially proposed by Allport (1954) in 1954 as the 'Contact hypothesis'. Drawing on studies from mixed and segregated neighbourhoods in the US, Allport concluded that under certain conditions, white people with frequent contact with black people experienced decreased racial prejudice. Despite being frequently construed as such, Intergroup contact theory is *not* a proposition of *frictionless* interactions of out-groups resulting in increased trust or social cohesion. Positive contact can potentially lead to these outcomes, but negative contact implies opposite effects (Pettigrew et al., 2011). Physical proximity increases the likelihood of contact, but whether that contact is positive (promoting understanding) or negative (invoking a threat perception) is not always clear (Pettigrew et al., 2011). In many empirical cases which often times feature migrants as an out-group, the contact conditions are not positive (Pettigrew et al., 2011). Migrants that flee poverty and seek work in a first-world country may have a very different collective set of common goals than the host society that was born into what they perceive as a very different status quo. Differing cultural norms between the host and migrant population can present a social challenge to migration (Collier, 2013), and this would not constitute a positive contact situation. Migrants might not speak the native language, presenting an obvious technical barrier to surpass, strengthening the 'otherness' perception of out-groups (Collier, 2013).

In the field of political science, contact theory has enjoyed increased attention since Robert Putnam proposed that, contrary to the consensus at the time, diversity decreases social cohesion in communities (Putnam, 2000, 2007). Putnam pointed out that while immigration as a source of diversity has a positive effect on society in the long run, its short-term effects can be largely negative. Collier (2013) echoes similar sentiments: in the long run, immigrants contribute to society and integrate, but in the short run, positive effects may be outweighed by social friction generated from the influx of diversity. The inclusion of migration into the Schelling model presented previously corroborates this view.

Contact and threat theory appear conflicting but focus on slightly different aspects of the same phenomenon. Threat theory is the study of precedents of prejudice towards outgroups, and contact theory is the study of the context in which different groups interact. The differences are described in further detail below.

Both theories have been empirically tested in social psychology and political science, resulting in hundreds of studies (for a meta-study of contact and threat theory, see Kaufmann and Goodwin (2016), for a meta-study on social capital and cohesion, see Portes and Vickstrom (2015)). To date, the relationship between diversity and social capital is unclear. The occurrence and strength of the relationship is dependent on the context it is placed in. Kaufmann and Goodwin (2016) find whether data is taken from national or sub-national level, will effect the resolution of information of groups and their behaviour and ultimately, the results of a study. Movement is easier to capture with higher resolution data. Tangential to these findings, Ariely (2014) notes the stark differences in operationalisation of social capital variables (particularly that of trust) that drive the differences in results between studies.

Empirical studies face difficulties in operationalising variables. Effects on the individual level, such as decreased prejudice as a result of positive contact with an out-group member, must not necessarily persist at the group level (Pettigrew and Tropp, 2006). However, much of the research that has spawned as a result of Putnam (2007)'s finding that diversity invoking threat perceptions on macro-level. This means that parts of the underlying theory cannot be captured. Threat theory emphasizes the threat-related antecedents of prejudice such as loss of identity though the presence of an outgroup that challenges the values on which the identity is founded. Contact theory by contrast focuses on the context of the contact (Stephan et al., 2009).

This chapter approaches inter-group tolerance from an agent-based perspective in order to understand the implications of migration as an introduction of diversity into an existing population. Previous research has employed agent-based models to introduce differing levels of tolerance (Hatna and Benenson, 2015b) and to explore the minority-majority relationships of different groups (Hatna and Benenson, 2012). In this chapter I build on the immigration model that was introduced in the previous chapter. I analyse a model in which I introduce a crucial innovation: tolerance in this model is adaptive; agents can alter their tolerance levels as they evaluate their surroundings. I use this model to investigate segregation outcomes under environmental conditions where migrants introduce new diversity into the existing population, and both groups have to adapt to the changed social environment. The adaptation of the model proposed by Hatna and Benenson (2015b) was chosen so that model outcomes of non-adaptive and adaptive agents can be compared more easily. The ways in which adaptive tolerance is implemented are based on the contact theory.

The remainder of the chapter is structured as follows. In the next section I describe my methods and my model. In Section 3.3 I present my results. Finally I conclude in Section 3.4.

3.2 Method

The method is once again ABM, using the Schelling model of the previous chapter. The main concepts of the Schelling model remain: agents have a preference for locations which are populated by agents of the own colour, and they move accordingly.

Preferences are quantified according to the threshold fraction of similarly-coloured agents in the neighbourhood that is required for an agent to be satisfied with its locale. Schelling (1971) showed that even a small preference to be near agents of the same colour gives rise to a large amount of segregation.

I use a similar framework, but introduce migration and adaptation of tolerance. I denote one of the colours — green — as representing natives, and the other — blue — as representing migrants. Migration is modelled by allowing the blue population to grow as new migrants arrive at particular times, and at particular locations, around which they cluster. Both groups of agents follow the same behavioural rules, which comprise a movement rule, and a tolerance adaptation rule. The former is similar to earlier Schelling models in which agents move over time relocating to their preferred neighbourhoods. The latter is an innovation of my particular model; when agents are exposed to the out-group their tolerance increases if they are currently satisfied with their environment, but otherwise it decreases.

In the next section I describe my model in precise detail. The model is analysed by simulating it very many times, recording and drawing free parameters randomly as described in Section 3.2.2. I analyse the model under five different immigration treatments, which are described in Section 3.2.3. I record the dependent-variables for each simulation run, as described in Section 3.2.4. In Section 3.3 I present a cross-sectional and time-series analysis of dependent and independent variables under each treatment. My analysis shows that there are very clear effects, which can be demonstrated by the use of simple descriptive statistics and scatter-plots. The source-code used for simulations is freely available under an open-source license (Urselmans, 2017a).

3.2.1 The model

The model baseline is very similar to the one described in the previous chapter, Section 2.2.2. To ease the comparison, most of the notation and variables are kept the same, any differences will be highlighted throughout the description of the model below. Again, a set of agents $A_t = \{a_1, \ldots, a_{n,t}\}$ are placed on a grid with a total of $N = 50 \times 50$ vertices V at time $t \in \mathbb{Z}$. Each agent a_i has a colour c_i, which is either blue ($c_i = B$), or green ($c_i = G$). Green agents, as previously, are the natives and they are randomly placed at the beginning of each simulation. Blue agents are migrants and arrive later on during migration waves. Agents cannot die or otherwise exit the grid. Fig 3.1 shows a visualisation of the now familiar patterns of the model: blue and green agents occupy one tile at most, and white space is empty space that agents can move to.

Figure 3.1: An example state of the simulation showing the colour c_i of each agent a_i. Blue squares are occupied by migrant agents and green squares by natives. White squares are empty cells. Both populations eventually form visible clusters.

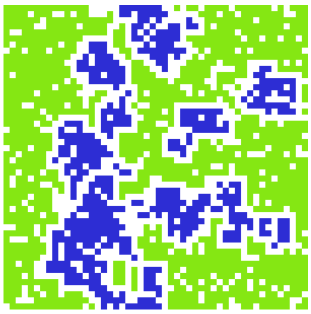

Agents can only see their local neighbourhood. Let $N_t(a_i)$ denote the neighbourhood of a_i at time t, which consists of the set of all other agents located on the lattice within a Euclidean distance of two nodes from a_i. The neighbourhood for a given agent thus consists of all other agents within its 5×5 Moore neighbourhood.

As implemented previously, agents have a preference to be with others of their own colour, but the implementation differs in this version of the model. Each agent a_i has a tolerance threshold $f_{i,t} \in [f_{min}, f_{max}]$ which determines the fraction of out-group members the agent tolerates in their immediate neighbourhood[8]. This fraction is continuous and differs from the discrete number of friends in Chapter 2.

[8]Despite the similarities to the model from 2, the mechanisms of tolerance deviate significantly and thus the F notation was dropped. Because tolerance is now continuous as it allows for slow changes, the scale of 24 discrete friends was no longer applicable.

The fraction of agents that are similar to agent a_i is given by

$$s_{i,t} = \frac{|\{a_j \in N_t(a_i) : c_i = c_j)\}|}{|N_t(a_i)|}. \tag{3.1}$$

Agents are satisfied with their neighbours if and only if the fraction of nearby similar agents meets their tolerance threshold. The utility of agent i at time t is again denoted by $u_{i,t} \in \{0, 1\}$, given by:

$$u_{i,t} = \begin{cases} 1 & : \quad s_{i,t} \geq f_{i,t} \\ 0 & : \quad s_{i,t} < f_{i,t} \end{cases}. \tag{3.2}$$

The movement rule algorithm for agents (3.1 below) incorporates the ratio s of in-group agents g and out-group agents d in an agent's neighbourhood.

Algorithm 3.1 Movement rule for agent a_i

```
L ← RandomlyChooseVacantSites(z)        ▷ choose |L| = w candidate locations
L⋆ ← {p_{i,t}}                           ▷ initialise satisfactory locations
for all l ∈ L do
    g ← |{a_j ∈ N_t(l) : c_i = c_j}|     ▷ number in-group agents in neighborhood
    d ← |{a_j ∈ N_t(l) : c_i ≠ c_j}|     ▷ number of out-group agents in neighborhood
    s ← g/g + d
    if d > 0 ∧ s ≥ f_i then
        L⋆ ← L⋆ ∪ {l}                    ▷ update satisfactory locations
    end if
end for
l⋆ ← ChooseOneAtRandom(L∗)
p_{i,t+1} = l⋆                           ▷ update location
```

Dissatisfied agents are given by $D_t = \{a_i \in A_t : u_{i,t} = 0\}$. Every tick or turn, each dissatisfied agent $a_i \in D_t$, who is currently located at $p_{i,t}$, randomly samples a number, z, of unoccupied locations L_i from the grid. They then randomly choose a new location from this subset for which the ratio of in-group to out-group agents meets their tolerance threshold, i.e. $\{l \in L_i \cup p_{i,t} : s_{i,t} \geq f_{i,t}\}$. If no satisfactory alternative locations are found, then the agent remains (unhappily) at its current location $p_{i,t}$. If an agent fails to find a new location, the agent's tolerance will remain unchanged. If the agent is still unhappy in the next time period, it will try and find a new location again.

As before, agents that are satisfied will have a small chance (probability 10^{-2} per tick) of relocating by randomly picking a location from z vacant locations without considering their utility. This aims to model movement that is not just due to tolerance of diversity, but for other reasons. The agents' reasoning, including their adaptive tolerance, is described in Algorithm 3.2 below:

Algorithm 3.2 Decision rule of agent a_i.

```
if |{a_j ∈ N_t(a_i) : c_j ≠ c_i}| > 0 then          ▷ at least one outgroup agent in neighbourhood?
    if u_{i,t} = 0 then                              ▷ agent is dissatisfied?
        Move agent                                   ▷ see Algorithm 3.1
        f_{i,t+1} ← f_{i,t}                          ▷ tolerance remains the same
    else
        f_{i,t+1} ← min(f_{i,t} + Δ_f, f_max)        ▷ increase tolerance by Δ_f
        p ← draw randomly from U(0,1)
        if p ≤ 0.01 then                             ▷ satisfied agents move with probability 0.01
            L ← RandomlyChooseVacantSites(z)
            p_{i,t+1} ← ChooseOneAtRandom(L)
        end if
    end if
else
    f_{i,t+1} ← max(f_{i,t} − Δ_f, f_min)            ▷ decrease tolerance by Δ_f
end if
```

Implementing the assumptions of contact theory, the model includes that tolerance of an agent adapts to its local environment. Positive contact with out-group agents leads to an increase in tolerance, in line with the theory (Allport, 1954). Positive contact is defined as sharing the neighbourhood with out-group members whilst the agent is happy. Accordingly, at each time period every satisfied agent $a_i \in A_t : u_{i,y} = 1$ who is exposed to at least one out-group agent in its environment increases its tolerance threshold by a constant term Δ_f, up to the maximum value f_{max}. Tolerance decreases by the same amount if an agent is surrounded by agents of the same colour, and is not in a contact situation: the reasoning is that lack of contact increases likelihood of stereotype-reliant views of the out-group (Pettigrew et al., 2011). In all other cases, tolerance remains the same:

$$f_{i,t+1} = \begin{cases} \min(f_{i,t} + \Delta_f, f_{max}) & : \quad u_{i,t} > 0 \wedge |\{a_j \in N_t(a_i) : c_j \neq c_i\}| > 0 \\ \max(f_{i,t} - \Delta_f, f_{min}) & : \quad |\{a_j \in N_t(a_i) : c_j \neq c_i\}| = 0 \\ f_{i,t} & : \quad u_{i,t} = 0 \wedge |\{a_j \in N_t(a_i) : c_j \neq c_i\}| > 0 \end{cases} . \quad (3.3)$$

Agents keep their tolerance level when they are unhappy and in a contact situation- this is because they will attempt to move away. Moving away *and* lowering tolerance would in effect constitute a double penalty of negative contact, as agents in a Schelling model will seek happiness not just through adaptation, but relocation. If agents were to lower their tolerance and move away, they would contribute to more segregation (by seeking a same-coloured neighbourhood meeting their needs) and to lower tolerance at once. However, it is possible to imagine why agents should carry out a double penalty:

Essentially, one has to decide whether it would be more realistic to implement the double penalty or not. Thinking about a real-life situation, a family of white people lives in a majority white neighbourhood. Then, black people start moving in. The ratio of white-to-black changes, and the white family decides they are not comfortable in this situation, and are going to move away. There are two potential causes: (1) the family is already intolerant, but doesn't know it yet. The exposure to black people merely reveals this low level of tolerance. (2) The family is somewhat tolerant, but the presence of black people *lowers* the existing tolerance, and this then leads to the unhappiness of the white family. In the first case, the family would not adjust their tolerance. In the second case, the family would indeed first lower their tolerance and then move away. It is thus possible to construct the rule in this different way, with potentially big impacts on the model outcomes. I have decided against the double penalty because a volume of research in social psychology, communication studies and discourse analysis on the denial of racism ("I have nothing against blacks, but..." in van Dijk (1992)) suggests that racism and prejudice more broadly are inherent and often unintended (Dovidio and Gaertner, 2010). Issues of inherent & systemic racism as well as the mixed results of anti-racism campaigns (Stephan et al., 2002) all suggest that a certain level of racism is present, and that this level is higher than people admit, even to themselves, and not just in one type of country (see for example Wasserman (2010) for a study in South Africa, Rojas-Sosa (2016) in Latin America, Andreouli et al. (2016) in the UK). There is also support (De França and Monteiro, 2013) for Dovidio's and Gaertner's theory of 'aversive racism', the idea that out-group prejudice is not exercised in situations in which the out-group is seen as positive, but that prejudice shows when situations are ambiguous (Dovidio and Gaertner, 2000).

Thus, it is likely that as underlying racist views are frequently denied, they are simply bought to the fore in a contact situation, rather than created.

A different approach would be to emphasise the volatility of racist thoughts and views based on framing (Branscombe et al., 2007), arguing that as people can change their minds quickly, they could decrease their tolerance before moving away. Moreover, a much-studied theme in intergroup relations is the use of legitimation of inequalities as a key driver perpetuating these very inequalities, including racism (Costa-Lopes et al., 2013). There is thus a material interest in keeping racist attitudes (or prejudice in general) up. I acknowledge the validity especially of the second argument. I argue that the act of moving in the real world constitutes a decision with many consequences, and is thus not taken lightly- a single situation framed in a negative way may not be enough to lead to such a drastic change- unless it simply exposes an intolerance that has been present, but had not been 'triggered' before. The need to keep racism up must not necessarily result in an increase in intolerance, it could also be interpreted as a refusal to decrease intolerance- something that is not discussed in contact theory and is not part of the model. To conclude, there are grounds on which the decision rule could include a double penalty, but there exists plenty of evidence that supports my model design of leaving it out.

The changes in tolerance in the model are tightly linked to states of happiness, and contact is not frictionless. The inclusion of the happiness-condition emulates the fact that integration is potentially a costly process. The costs themselves are not modelled, but rather, the happiness condition acts as a proxy for willingness to pay these costs. A happy person is more willing to engage than an unhappy person. Happiness, I recall, is purely a representation of whether the neighbourhood is satisfactory.

Depending on the experimental treatment (see Section 3.2.3 below), the population of agents can grow as new migrants arrive. The number of native agents is always constant, new migrants arrive in discrete waves of migration up until the final point of arrival, $t_{mig} = 1000$. The population dynamics are specified in terms of:

1. the final population density — $PopDen$ — which is a parameter that specifies the fraction of occupied sites after all migration event have occurred (at $t = t_{mig}$);

2. the native share of the population — $NatShare$ — which specifies the ratio of natives to migrants at the end of the simulation; and

3. the number of waves of migration — E — which specifies how many migration events occur.

Migration is carried out using the same influx mechanism *Cluster* introduced in the previous chapter (see Section 2.2.3. Migrants will target high-density areas irrespective of whether migrants or natives populate those areas. At the beginning of the simulation a total of N_G native agents are placed randomly onto the lattice:

$$N_{max} = \text{round}(PopDen \times N) \tag{3.4}$$

$$N_G = \text{round}(NatShare \times N_{max}), \tag{3.5}$$

and for the control where there is no immigration ($E = 0$) a total of N_B migrant agents are also placed randomly at the start of the simulation, making them in effect blue natives:

$$N_B = \text{round}((1 - NatShare) \times N_{max}) \tag{3.6}$$

In treatments where migration is dynamic ($E > 0$), there are no migrant agents on the lattice at the beginning of the simulation. Rather, the first wave of migration occurs at time $0.05 \times t_{mig}$, and the subsequent migration waves occur at evenly spaced intervals of $0.9 \times t_{mig}/E$ ticks. During each wave of migration an additional number

$$\Delta_B = \text{round}(N_B/E) \tag{3.7}$$

of migrant agents are placed onto the lattice, clustering around a focal vacant location p_B. The additional migration sites for the new arrivals are chosen by iteratively finding the best neighbour of the chosen focal location p_B; sites are ranked firstly according to the highest number of surrounding new migrants, and secondly according to their local population density within their neighbourhood. The placement algorithm is summarised in the pseudo-code given by algorithms 3.3 and 3.4.

Algorithm 3.3 Choose locations for migrant agents during migration waves

```
function PlaceMigrants(p_B, Δ_B)            ▷ Place Δ_B migrant agents around location p_B
    P_B ← {p_B}                               ▷ Initialise the set of locations for immigration
    while |P_B| < Δ_B do                      ▷ More migrants to place?
        p_B ← BestNeighbour(p_B, P_B)         ▷ Find the best neighbouring location
        P_B ← P_B ∪ p_B                       ▷ add it to the result set
    end while
    return P_B
end function
```

Algorithm 3.4 Find the neighbouring site with the greatest population density

```
function BestNeighbour(p_B, P_B)   ▷ Best neighbour of p_B excluding locations P_B
    if |N(p_B) − P_B − {a_i : p_i ∈ N(p_B)}| > 0 then   ▷ Vacant sites not already chosen?
        P* ← {}                                          ▷ Initialise best locations
        d* ← −∞                                          ▷ Initialise best density
        for all p ∈ N(p_B) − P_B − {p_i : a_i ∈ A_t} do  ▷ All vacant unchosen neighbours
            d ← |{a_i : p_i ∈ N(p)}|/|N(p)|             ▷ Calculate local population density
            if d > d* then
                d* ← d
                P* ← P* ∪ (p, d*)
            end if
        end for
        return ChooseOneAtRandom({p : (p, d) ∈ P* ∧ d = d*})
    else
        p ← ChooseOneAtRandom(N(p_B))
        return BestNeighbour(p, P_B)
    end if
end function
```

3.2.2 Initial conditions

The majority of parameters governing the initial conditions of the model are randomly varied between simulation runs in order to test the robustness of the model. I also record these values so that they can be used as independent variables in order to ascertain any effects. These are described in turn below, and summarised in Table 3.1 (the remaining constant parameters are summarised in Table 3.3, and the state variables in Table 3.4).

Tolerance distribution

When agent a_i arrives at the simulation its initial tolerance $f_{i,0}$ is drawn *i.i.d.* from a uniform distribution $f_{i,0} \sim U(f_{min}, f_{max})$. After the initialization, the agent adapts their tolerance according to Equation 3.3 as summarised in Algorithm 3.2. For all simulations in this chapter I set $f_{min} = 0.05$ and $f_{max} = 0.95$. These limits prevent agents from 'locking in' at the extreme values of tolerance; with these constraints agents will always be able to tolerate one out-group member in their neighbourhood (without these constraints, any agent reaching full tolerance or intolerance would never readjust again, since just one out-group member would be above the tolerance threshold). Due to the conversion of discrete numbers of friends between 0 to 24 (in the neighbourhood) to a continuous scale of tolerance between 0 and 1, multiple values of tolerance correspond to one discrete friend, as agents can only be whole. For instance, tolerance levels between 0 and 0.04 all correspond to 0 friends.

Rate of change of tolerance

The rate of change of tolerance Δ_f is the increment used when agents adapt their tolerance (see Equation 3.3). At the beginning of each simulation it is drawn randomly $\Delta_f \sim U(10^{-5}, 10^{-3})$ and remains constant throughout the simulation. The low values of Δ_f are chosen to reflect the fact that attitudes, on average, change only slowly. Because the time-scale of the model is arbitrary (one tick could mean a day, a month, a year, an electoral cycle), the rate of change of tolerance in relation to the number of ticks can determine on what scale the model is interpreted. A Δ_f of 0.001 means that an agent has to increase tolerance for 40 subsequent ticks before one additional out-group member will be tolerated. This is a very slow pace, designed to model a society over several decades. Adjusting the rate of change of tolerance can represent faster-paced time-scales of months or a couple of years.

Final population density

The final population density $PopDen$ determines the fraction of occupied sites after all waves of immigration have occurred ($|A_{t_{mig}}|/N$). From thereon, the number of agents is constant. At the beginning of each simulation this parameter is randomly drawn from a uniform distribution $\sim U(0.75, 0.98)$.

Schelling models are typically assumed to run under conditions of high density (Hatna and Benenson, 2012), which is why the minimum is still 3/4 of the map covered. Density also acts as a proxy for freedom of choice. Higher density results in less freedom of choosing better areas.

Final native share of the population

The final native share of the population $NatShare$ determines the ratio of natives to migrants after all waves of immigration have occurred. This parameter is initialised randomly by drawing from a uniform $\sim U(0.02, 0.98)$ at the beginning of each simulation. In treatments without migration ($E = 0$) it determines the fraction of natives in the initial population, which thereafter remains fixed. In treatments with migration ($E > 0$), it determines the number of migrants added in each wave (see equations 3.5, 3.6 and 3.7), which in turn determines the final fraction of natives in the population.

Considering the extremes of this parameter, when $NatShare = 0.02$, the world would fill up with migrants until migrants constitute 98% of the population, and natives constitute 2%. Whilst national-level migration does not lead to migrants outnumbering natives, the reasoning is that on smaller geographical areas, this majority-minority flipping can indeed occur. Because segregation is mediated by how society is made up, how big minorities are and how they are distributed, the ratio seeks to test in how far, if at all, different minority-majority relationships influence segregation behaviour and tolerance levels. The traditional Schelling model has usually assumed an even split, an assumption that is not theoretically useful in the context of migration and attitudes towards diversity.

Considered tiles to move

The parameter z specifies the number of vacant locations than each agent considers when moving. At the beginning of each simulation it is initialised randomly by drawing from a discrete uniform distribution $z \sim U(25, 125)$. The minimum and maximum of this distribution correspond to 1% and 5% respectively of the size of the entire lattice. In the previous model in Chapter 2, this parameter was fixed as it did not affect any model outcomes. Because there are a number of significant changes in this version, the parameter was re-introduced to see whether it would have an effect this time round.

3.2.3 Immigration treatments

Consistent with the immigration implementation of the previous model, there are five different experimental treatments for immigration of blue agents into the model, which correspond to five different values of the parameter E. These are summarised in Table 3.5. Immigration waves arrive within the first 1,000 ticks of $t_{max} = 20,000$. This allows agents to adjust their behaviour for a prolonged period after the last migration wave has occurred.

The first treatment $E = 0$ is a control condition with no immigration taking place. In this condition, both natives and migrants are initialised at the start of the simulation and there is no increase in the size of the migrant group over time. The four remaining treatments all feature immigration at different rates, aiming to simulate one-off large influxes of migrants as well as a "trickle-down" scenario in which few migrants arrive at one time, but do so for a sustained period of time. The precise dynamics are described by equations 3.4 to 3.7 in the previous section. For a more in-depth description with visualisation, see Section 2.2.3 in Chapter 2.

3.2.4 Dependent variables

Dependent variables are sampled on every simulation run every 10 time steps, allowing for both cross-sectional and time-series analysis. These variables are described below, and summarised in Table 3.6.

Firstly, the Moran's Index of spatial autocorrelation is recorded for (1) colour segregation:

$$M_t^c = \frac{|A_t|}{\sum_{(i,j)\in A_t^2} w_{i,j}} \frac{\sum_{(i,j)\in A_t^2} w_{i,j}(c_i - \bar{c}_t)(c_j - \bar{c}_t)}{\sum_{i\in A_t}(c_i - \bar{c}_t)^2} \tag{3.8}$$

where the mean colour is $\bar{c}_t = \sum_{i\in A_t} c_i/|A_t|$, and $w_{i,j} = 1$ if and only if agents a_i

and a_j are immediately adjacent on the grid, including diagonals, otherwise $w_{i,j} = 0$, and (2) segregation of *tolerance*, M_t^f, by substituting f in place of c in equation 3.8:

$$M_t^f = \frac{|A_t|}{\sum_{(i,j)\in A_t^2} w_{i,j}} \frac{\sum_{(i,j)\in A_t^2} w_{i,j}(f_i - \bar{f}_t)(f_j - \bar{f}_t)}{\sum_{i\in A_t}(f_i - \bar{f}_t)^2} \tag{3.9}$$

Because happiness is now in part a function of adaptive tolerance, it is no longer a dependent variable. Instead, I focus on the tolerance of agents: I record the first four moments of the tolerance distribution across the population ($\bar{f}_t$, $\sigma^2_{f_t}$, γ_{f_t} and κ_{f_t}), and subdivide this into tolerance of migrants:

$$\bar{f}_{B_t} = \sum_{a_i \in B_t} f_i/|B_t| \tag{3.10}$$

and tolerance of natives:

$$\bar{f}_{G_t} = \sum_{a_i \in G_t} f_i/|G_t| \tag{3.11}$$

where B_t is the migrant population $\{a_i \in A_t : c_i = B\}$, and G_t is the native population $\{a_i \in A_t : c_i = G\}$.

Finally, to help identify whether the tolerance distribution is bimodal, I record the bimodality coefficient (Pfister et al., 2013) of the tolerance distribution:

$$\beta_{f_t} = \frac{\gamma^2_{f_t} + 1}{\kappa_{f_t}} \tag{3.12}$$

As a notational convention, I refer to the final value of an independent variable at $t = t_{max}$ by omitting the time subscript from all of the above.

Table 3.1: Independent variables

Parameter	*Distribution*	*Description*
$NatShare$	$\sim U(0.02, 0.98)$	Fraction of natives
$PopDen$	$\sim U(0.75, 0.98)$	Final population density
Δ_f	$\sim U(10^{-5}, 10^{-3})$	Rate of change of tolerance
z	$\sim U(25, 125)$	No. of considered locations when moving

Table 3.2: Restricted ranges of independent variables for which the model converges within $t \leq t_{max}$

Parameter	*Range*
$NatShare$	$0.2 \leq NatShare \leq 0.8$
Δ_f	$5 \times 10^{-4} \leq \Delta_f < 0.001$

Table 3.3: Constants

Constant	*Description*
$t_{mig} = 1,000$	Time until final migration
$t_{max} = 20,000$	Maximum number of ticks per simulation
$N = 50 \times 50$	Size of lattice
$f_{min} = 0.05$	Minimum tolerance
$f_{max} = 0.95$	Maximum tolerance

Table 3.4: State variables

Variable	*Description*
A_t	The population of agents
$c_{i,t}$	Colour of agent a_i
$f_{i,t}$	Tolerance of agent a_i
$u_{i,t}$	Utility of agent a_i
$p_{i,t}$	Position of agent a_i
$N_t(a_i)$	The set of agents that are neighbours of agent a_i
$N(p)$	The set of locations in the neighbourhood of location p

Table 3.5: Treatment conditions

Treatment	$E = 0$	$E = 1$	$E = 4$	$E = 15$	$E = 100$
Migration	No	Yes	Yes	Yes	Yes
Number of waves	-	1	4	15	100

Table 3.6: Dependent variables

Variable	*Description*
M_t^c	Segregation of colour at time t (equation 3.8)
M_t^f	Segregation of tolerance at time t (equation 3.9)
$\bar{f}_{B_t}$	Tolerance of migrants at time t (equation 3.10)
$\bar{f}_{G_t}$	Tolerance of natives at time t (equation 3.11)
β_f	Bimodality of tolerance at the end of the simulation (equation 3.12)
M^c	Segregation of colour at the end of the simulation
M^f	Segregation of tolerance at the end of the simulation
$\bar{f}_G$	Tolerance of natives at the end of the simulation
$\bar{f}_B$	Tolerance of migrants at the end of the simulation

3.3 Results

The model was analysed through simulation and empirical methods. For each treatment in Table 3.5 I executed 7,000 independent realisations of the model, drawing free parameters from the distributions specified in Table 3.1. Each realisation was run for a total of $t_{max} = 20,000$ simulation ticks. This resulted in a total of $5 \times 7,000 = 35,000$ cross-sectional samples of each of the dependent variables (Table 3.6). During each simulation I also sample all dependent-variables every 10 ticks, resulting in a total of $35,000 \times (20,000/10) = 7 \times 10^7$ time-series samples. In the following I first give an overview of the qualitative properties of a single typical simulation run before analysing the aggregate data across simulation runs. The model is updated sequentially, every tick. Agents act in sequential order; that order is shuffled every tick to avoid tactical advantages that might result from a pre-determined sequence.

3.3.1 A typical simulation run

Fig 3.2 shows a visualisation of a typical simulation run to demonstrate the clustering of agents. The top row shows a progression of migration ($E = 4$) at a starting density of 37%, filling up with migrants up until a 75% density so that both groups are equal in size. The blue and green agents move around, empty space is white.

The bottom row shows the corresponding tolerance heat-map. Light colours denote tolerant agents, dark colours denote intolerant agents, and the grey areas are vacant tiles. The period from $t = 90$ (pre-migration) to $t = 100$ (post-migration) is marked by a significant changes in tolerance levels. The native fraction of the population that is not in vicinity of migrants are uniformally hostile, whereas the newly arrived migrants have a large variance in their tolerance levels, which are randomly drawn from a uniform distribution upon entering the map. At this stage, the map is sparsely populated and the clusters of migrant and natives have visible buffer-zones between which make inter-group contact less likely.

Figure 3.2: States of the simulation at different times. The top row shows the colours c_i of each agent. The bottom row shows the corresponding tolerance heatmap.

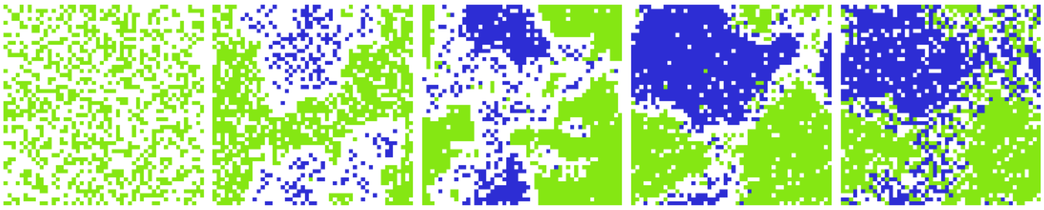

$t = 90$ $t = 100$ $t = 300$ $t = 790$ $t = 990$

(a)The colour c_i of each agent at $t = 90, 100, 300, 790, 990$ showing the first arrival of migrants

Native Migrant Empty

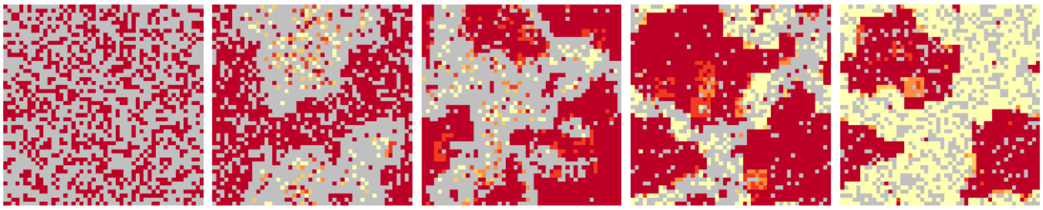

$t = 90$ $t = 100$ $t = 300$ $t = 790$ $t = 990$

(b) The tolerance f_i of each agent as a heat-map for the corresponding states above.

f_i : $< .2$ $> .2 < .4$ $> .4 < .6$ $> .6 < .8$ $> .8$

As more migrants arrive, I start to observe the effects of inter-group contact. Natives exposed to migrants react either by increasing their tolerance, resulting in the lighter colours visible at $t = 300$, or by moving. Once all of the migrants arrive, at $t = 790$, there are two pronounced clusters of agents (one of which wraps around the grid). On the tolerance heat-map, I see corresponding clusters of tolerance, with highly-intolerant agents surrounded by highly-tolerant agents. At this stage, most of the population has either very high or very low levels of tolerance, but as long as there is enough empty space to form a buffer zones between the clusters, the majority of agents are still highly intolerant, because only very few contact situations arise; unhappy agents relocate before adapting their tolerance levels.

This situation changes as space becomes more scarce. The final two states shown in Fig 3.2 illustrate a rapid phase-transition from a mainly intolerant society into a bimodal society of two equally-large fractions of highly-tolerant and highly-intolerant agents. The vacant buffer zones are now populated with tolerant agents of both colours who have relocated from the periphery of their respective clusters, thus forming a new zone of highly-tolerant agents. This tolerant zone expands as more out-group members mix in these high-tolerance areas, which in turn influence agents on the periphery of the intolerant clusters through positive contact, which results in a rapid erosion of their size.

This process continues until the intolerant clusters are completely surrounded by tolerant agents, who are satisfied and therefore static. This provides a dense, rigid substrate which restricts the movement of agents on the periphery of the intolerant clusters, who provide a protective membrane shielding the inner-core from further out-group contact. The periphery itself is highly dynamic; because agents on the periphery are unhappy they relocate, but their range of movement is restricted to locations within or near the cluster. However, inside these clusters, agents are intolerant but satisfied, as they are surrounded by in-group members, and therefore they remain static. Thus the entire cluster of intolerant agents achieves a relatively stable configuration, and persists over time.

The smaller orange-coloured clusters are unstable pockets of medium tolerance which appear throughout the simulation, but rapidly disappear again as they consist entirely of satisfied agents who either become more tolerant through out-group contact, or become isolated and intolerant.

Thus movement and adaptive tolerance interact leading to an emergent shield-and-buffer dynamic that polarises the population, causing it to self-assort along the tolerance axis with agents being either extremely tolerant or intolerant. Although in this section I have only discussed a single simulation run, in subsequent sections I show empirically that the model results in bimodal tolerance for many different initial conditions, and despite Monte-Carlo variance.

3.3.2 Model convergence

As discussed in the previous section, the model exhibits subtle dynamics, and the parameter range for the rate of change of tolerance Δ_f has been very low on purpose.

Therefore it is important to establish whether the key dependent variables stabilise within the finite time period $t \leq t_{max}$. I test the convergence of each independent realisation of the model individually by analysing the final $n = 250$ values V_f of the time-series of each independent variable V sampled at intervals of 10 ticks.

The criteria I use to test convergence of each variable are:

(i) if the variance of the final sample is extremely small $\sigma^2_{V_f} < 10^{-20}$; or

(ii) the standard deviation is small compared to the overall range $\sigma_{V_f} < [\max(V) - \min(V)] \times 10^{-4}$; or

(iii) if V_f is stationary under an augmented Dickey-Fuller test (Elliott et al., 1996). This is established by estimating the model $\Delta V_t = \alpha + \beta t + \gamma V_{t-1} + \delta_1 \Delta V_{t-1} + \ldots + \delta_{\lambda-1} \Delta V_{t-\lambda+1} + \epsilon_t$ where the lag order λ is chosen using the Akaike information criterion, and accepting the time-series as convergent *i.i.f.* if the value of the test statistic $\hat{\gamma}/SE(\hat{\gamma})$ is less than the critical value for $p = 0.05$.

These criteria were chosen because they allow to test not only for cases where the model reaches a static steady state in which values of dependent variables are constant over time, but also stochastic steady states in which the time-series is stationary; i.e. the *moments*, such as the mean and variance, are constant, despite the fact that the dependent variable has non-zero rate of change.

Using the above criteria, I analyse the time-series of the tolerance of migrants, the tolerance of natives, and the segregation of colour; i.e. $V \in \{\bar{f}_{B_t}, \bar{f}_{G_t}, M^c_t\}$. I record that the model has converged for a given realisation *i.i.f.* all three variables converge in the final period.

Over the entire range of parameters only 86% of simulation runs reached (stochastic) steady state within t_{max} ticks. I was able to identify which independent variables contributed to the convergence of the model through a correlation analysis, which identified $NatShare$ and Δ_f as the most promising explanatory variables. I binned

Figure 3.3: Proportion of simulations run that converge by independent variable. Each variable X is binned into intervals of size 0.05, and then I count the fraction of independent simulation runs which converge within each bin.

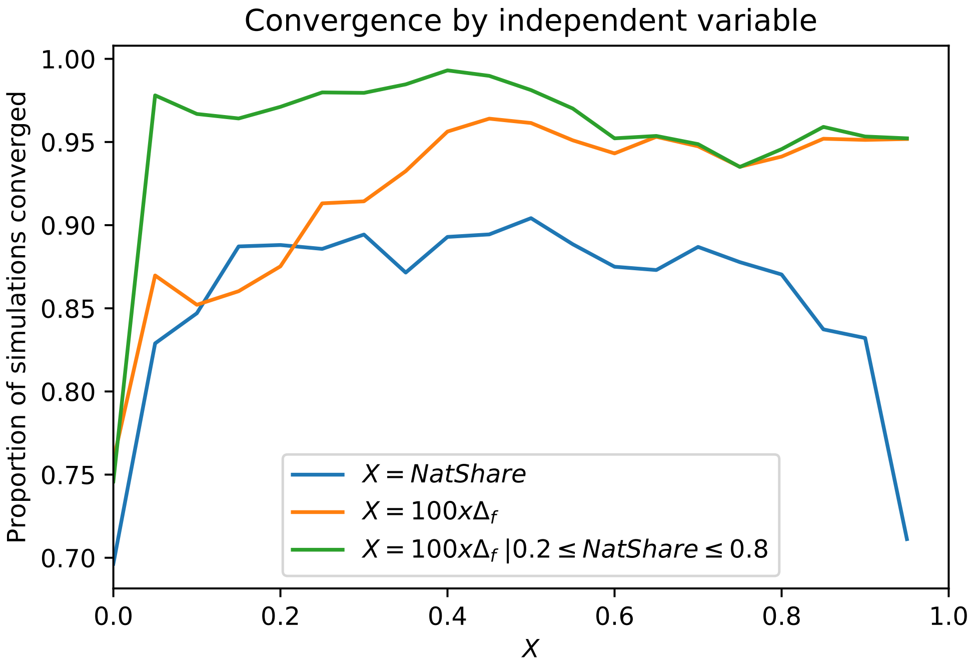

$NatShare$ and $\Delta_f \times 100$ into bins of size 0.05, and plotted the proportion of simulation runs that converged within each bin (Fig. 3.3).

The highest failure rate occurs for extreme values of $NatShare$ and for small values of $\Delta_f < 0.0005$; for very small values of Δ_f agents adapt very slowly, and the model fails to reach equilibrium within $t \leq t_{max}$, whereas for extreme values of $NatShare$, the respective minorities are likely too small to form sustainable clusters, thus breaking up and forming again. In fact, both variables interact; if I do not control for $NatShare$ then the convergence rate increases asymptotically, but slowly, with Δ_f. However, if I control for extreme values of $NatShare$ then provided that $\Delta_f \geq 0.0005$ I obtain a fairly consistent convergence rates of $\sim 97\%$.

Based on the convergence analysis, I restrict the ranges of the $NatShare$ and Δ_f parameters used in the remainder of the chapter. All results in subsequent sections use the ranges $0.2 \leq$NatShare≤ 0.8, and $\Delta_f \geq 0.0005$, for which the vast majority (97%) of simulation runs converge. These ranges are summarised in Table 3.2.

3.3.3 Time-series analysis

The tolerance of natives $\bar{f}_G$ (henceforth: 'native tolerance') and migrants $\bar{f}_B$ (henceforth: 'migrant tolerance') over time is plotted in Fig 3.4). An important variable affecting tolerance behaviour is $NatShare$, which determines the final size of the native share of the population. Fig 3.4 shows tolerance levels over time by the high and low values of native share, grouped into $NatShare \leq 0.3$ and $NatShare \leq 0.7$ respectively.

In the case of low native share, under no migration ($E = 0$) conditions, natives and migrants behave similarly (graphs Fig 3.4a and Fig 3.4f). The majority group is slightly less tolerant. The differences are visible at first, but nearly converge after $t = 2500$. The differences are very small, but still significant.

The case of $E = 1$ stands out from the rest. When $NatShare$ is low (Fig 3.4b), natives never recover fully from the initial shock of migration. As with all other cases, native tolerance drops sharply and quickly recovers, but not exceeding $\bar{f}_G = 0.7$ for the remaining time. The reverse scenario of high $NatShare$ (Fig 3.4g) is visibly different: natives recover and reach near-total tolerance, along with migrants. These values are slightly higher than those observed in Fig 3.4f, with no migration.

The reason that the $E = 1$ treatment is so different from the rest is that in the case of low $NatShare$, the proportion of natives to migrants is flipped immediately: native agents who constitute 20% of the final population make up 100% of the pre-migration population. Because of the large number of migrants coming in, many migrants are immediately exposed to a majority of migrants, presenting a shock not just to parts of the population, but to most of it. As no migrants had been present before, natives have not clustered into groups which could at least 'shield' the inner part from the effects of the one-off migration (this is because segregation behaviour is triggered by unhappiness rather than absence of out-group members, see section **??**). Because natives suddenly find themselves in a minority and most of the agents become unhappy, they will move into areas with fewer migrants. The sudden influx

Figure 3.4: Time series of native tolerance $\bar{f}_G$ (green) and migrant tolerance $\bar{f}_B$ (blue), by treatment, filtered by extreme values of the native-share initial condition ($0.2 \leq NatShare \geq 0.7$). The error bars show the 95% confidence interval for the mean of the $\bar{f}$ values across independent simulation runs and the range of Δf that leads to convergence. The dashed line at t_{mig} marks the end of migration waves.

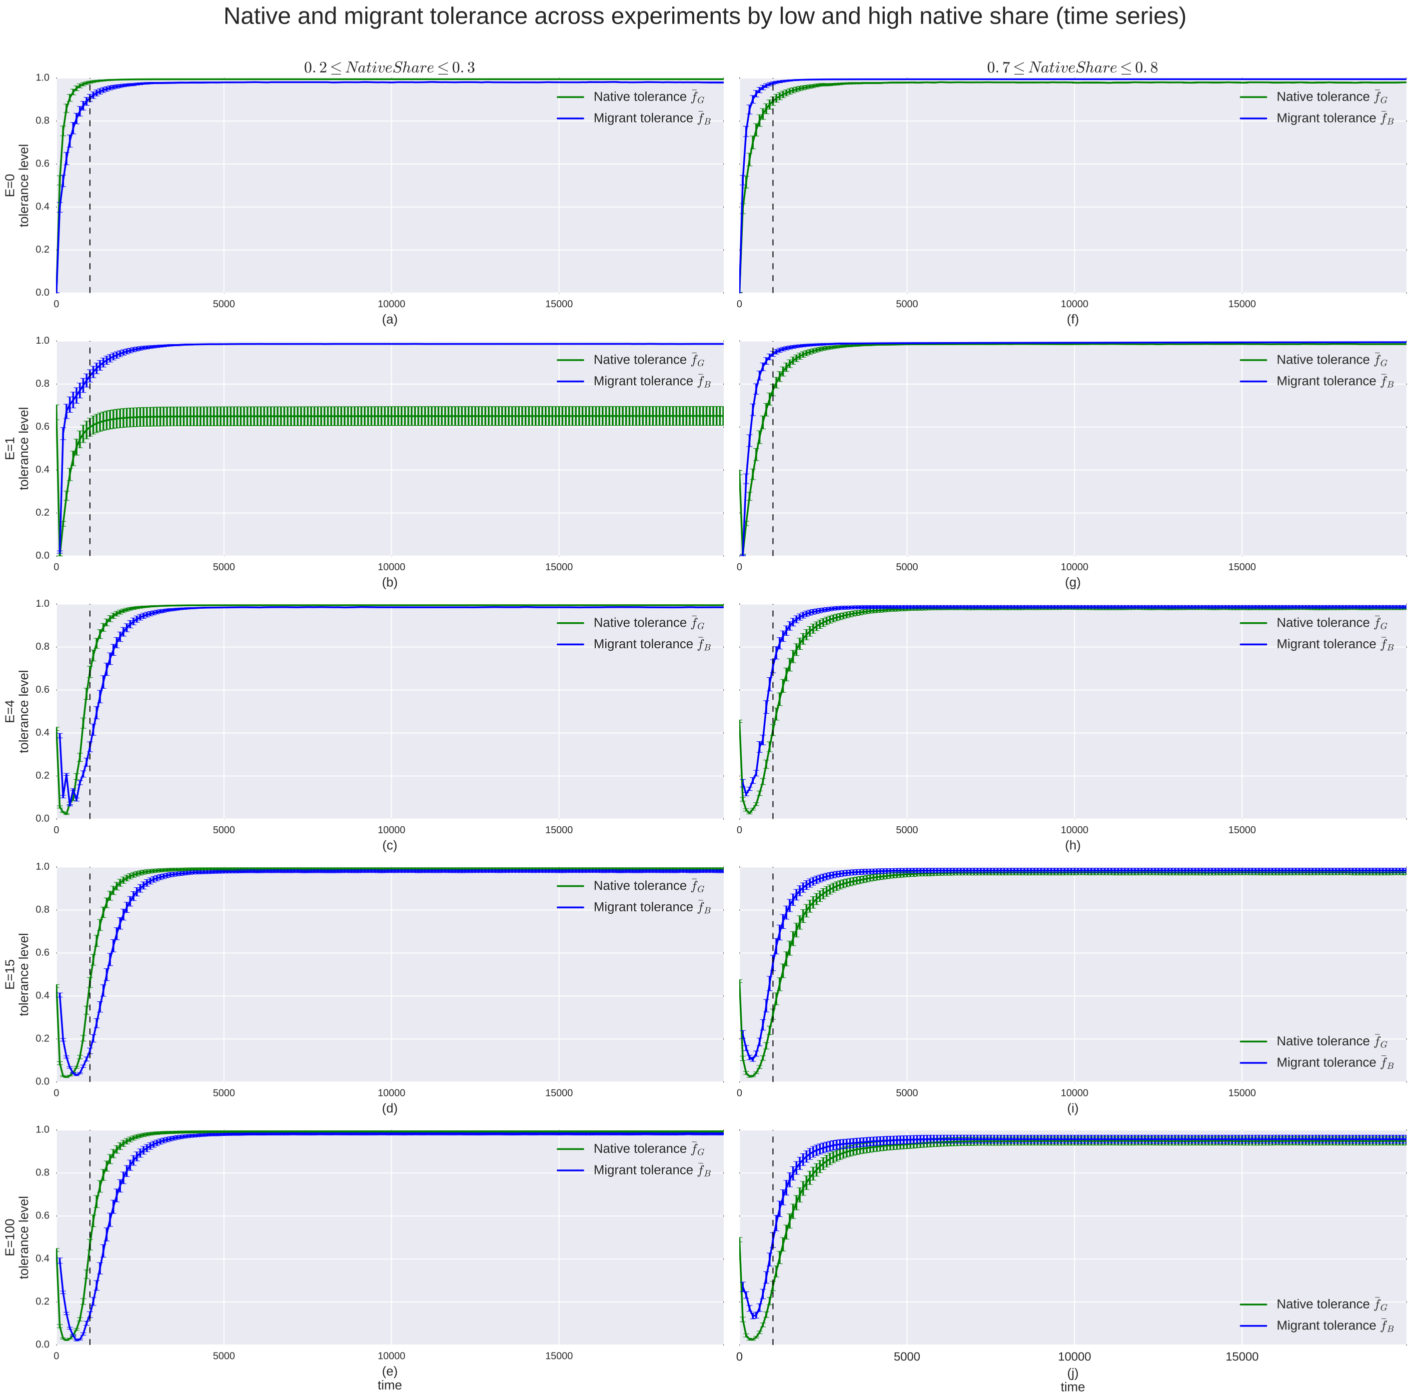

has prevented even moderate natives to adapt to their changing neighbourhood, and as they start to segregate away, not enough tolerant natives remain to become happy and increase their tolerance levels. The sudden minority is an important element of this behaviour. As the high $NatShare$ situation shows (Fig 3.4g), a one-off influx does not reduce long-term native tolerance when natives are in a majority even after the large migration wave. In contrast, natives in $E = 0$ treatments are exposed to migrants from the very beginning: they experience unhappiness and will segregate to improve it, creating pockets of happy and intolerant natives surrounded by tolerant natives that are exposed to migrants and shield the intolerant parts of their population group from exposure. This process is not achieved in $E = 1$ treatments, where previous absence of migrants has meant that native tolerance had started to decline.

At $E = 4$ (Fig 3.4c, Fig 3.4h), the shocks from each migration wave is visible. When $NatShare$ is low (Fig 3.4c), migrant tolerance is low when migration is still occurring. The increases in tolerance at each migration wave is due to the random initialisation of tolerance for new migrants. Effectively, each wave presents an opportunity to tip the tolerance balance within the population. This is not achieved until the fourth and last wave of migration has arrived at $t = 700$. Natives are visibly affected by the influx of migrants. During the first two migration waves, native tolerance $\bar{f}_G$ drops to near-zero, before increasing sharply to above 0.6 at $t = 1000$. Beyond this point, both natives and migrants transition to a majority tolerant society of $\bar{f}_G > 0.98$ by $t = 5000$.

For $E = 15$ (Fig 3.4d,i) and $E = 100$ (Fig 3.4e,j) the overall pattern is very similar. The migration waves are no longer as visible on the graphs, as the size of waves is not large enough to upset the overall population. Both natives and migrants will first experience a drop to low levels of tolerance, and recover quickly as more waves arrive, reaching their peak. Higher numbers of migration waves increase the time required to reach convergence of peak tolerance when $NatShare$ is high (Fig

3.4i,j). When $E = 4$, the peak is reached by $t = 5000$. When $E = 100$, this requires an additional 2,500 ticks. This means that for longer periods of time, the population groups are not as tolerant. The higher convergence times are also visible in Fig 3.3.

A notable difference between natives and migrants in the $E = 15$ and $E = 100$ cases is that natives will always drop their tolerance to near-zero, regardless of their population share. Migrants mirror this pattern only when they are in the minority (Fig 3.4d,e). When migrants are a majority (Fig 3.4i,j) their lowest tolerance is above 0.1. If natives and migrants behaved the same way, the graphs on the left should be mirrored by the graphs on the right.

Lastly, when $E = 100$ and $NatShare$ is high (Fig 3.4j), the maximum tolerance levels are not as high compared to cases with fewer migration waves. The variance increases as the number of waves increases, suggesting a less settled pattern of tolerance.

3.3.3.1 Non-convergent cases

The previous section highlighted that not all cases of the model converge, and that two parameters, namely $NatShare$ and Δ_f, had to be restricted to ensure model convergence. However it can be useful to investigate what happens in the case of non-convergence. Even though the end-state results cannot be supported, the behaviour of agents over time in cases of non-convergence can shed light on the model dynamics.

Figure 3.5 plots the tolerance of natives and migrants over time, just like Figure 3.4. In this case, the low range of $NatShare$ is between 0.01 and 0.1, and the high range between 0.8 and 1; the two ranges that lead to non-convergence. On first glance, the non-convergence occurs because the tolerance of each population group is usually still increasing by the time the simulation finishes. I note here that the model was also run at different lengths of up to $t_{max} = 200,000$ ticks, i.e. ten times longer than the current set of results. None of these repetitions led to convergence-instead, it would slow down, but still tend upwards at the final stage. Comparing the outcomes of Figure 3.4 and 3.5, tolerance is generally much lower in the non-convergent cases. In the convergent cases, total tolerance is reached quickly after the final migration wave has occurred. This is not the case here: Native tolerance, which is usually near-1 and only deviates from this pattern when $E = 1$ and $NatShare$ is low, is now often lower, reaches full tolerance when $NatShare$ is low but also when $E = 1$, reversing the pattern.

Notable is also that the high levels of tolerance do seem convergent for $E = 1$. The error bars show very low variation and the non-convergence in this case is largely driven by the erratic migrant tolerance (Fig 3.5b). Given that natives are in such a small minority in these cases, they are too few to find each other to form a coherent group, and adjust by increasing their tolerance towards the migrant group, which constitutes 90%+ of the population. When $NatShare$ is slightly higher, as is the case in Figure 3.4, it is enough to provide 'pockets of intolerance' which prevent natives from lowering their thresholds and thus remaining unhappy after the large shock of the one-off migration. Migrants are quite intolerant in the case of low $NatShare$, seen in Fig. 3.5a,b,c,d,e. Population groups that outnumber migrants to such a large extend simply can't provide sufficient contact opportunities for agents to increase their tolerance levels. The rate of change of arrival does not alter this pattern, however tolerance is higher overall when migrants arrive later on, as opposed to the control experiment of no-migration. Migrant happiness slumps after reaching a high population share (this pattern occurs in the non-adaptive model as well, see Section 2.3) and starts increasing again, very slowly and has not settled by the time t_{max} is reached.

When $NatShare$ is high (right hand side of Figure 3.5), both natives and migrants have moderate levels of tolerance and do not show the near-maximum levels that they show when the model converges. What is interesting here is that the high end of $NatShare$ does not mirror the low end of the range: now migrants are the

Figure 3.5: Time series of native tolerance $\bar{f}_G$ (green) and migrant tolerance $\bar{f}_B$ (blue), by treatment, non-convergent cases only. Native share is $(0.01 \geq NatShare \leq 0.1)$ and $(0.8 \geq NatShare < 1)$. The error bars show the 95% confidence interval for the mean of the $\bar{f}$ values across independent simulation runs. The dashed line at t_{mig} marks the end of migration waves.

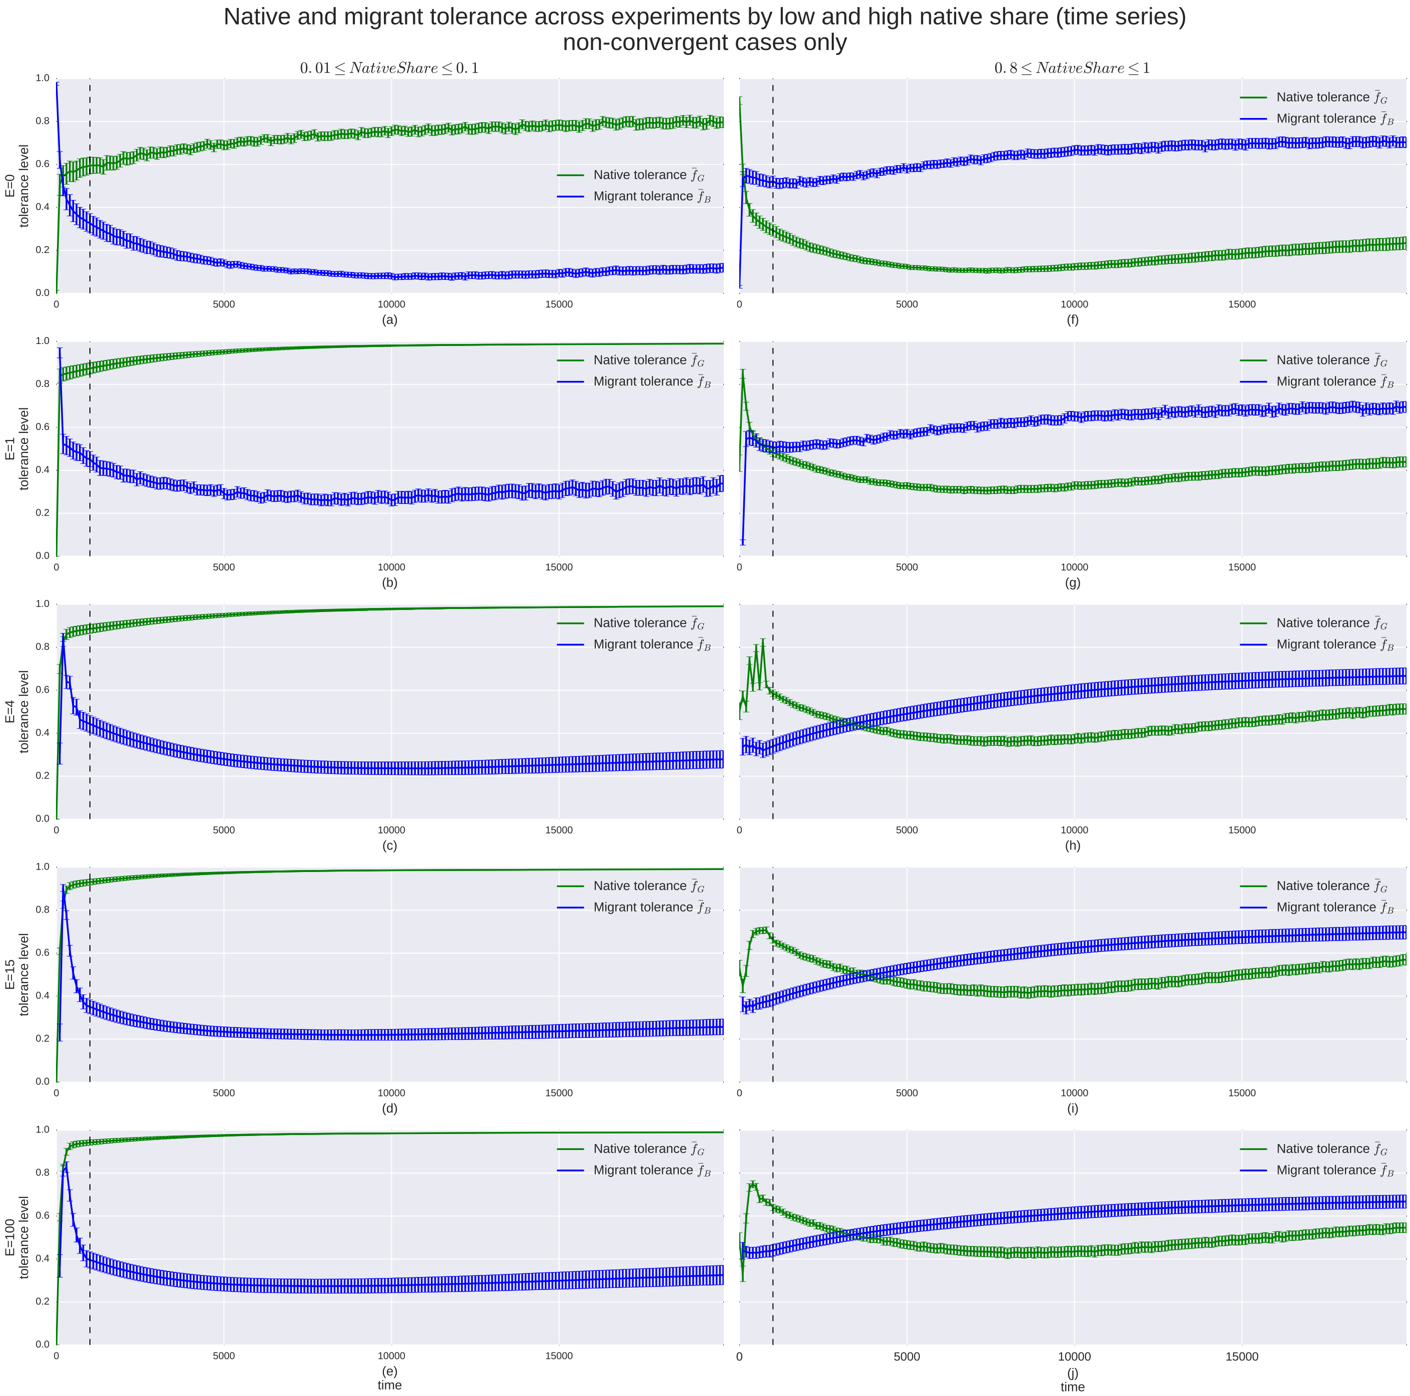

minority, yet they do not reach stages of full tolerance. Instead, migrant and native tolerance first flips (natives start off more tolerant and become the less tolerant population group) and then continues to increase, natives slowly catching up to migrant tolerant- but again, as the end-state of the model does not converge, there is no certainty that the trend continues upwards.

The non-convergent cases of the model also show a very different outcome of tolerance- it is not that the convergent cases show the same patterns with the only difference being that they reach convergence before the simulations terminate- the short-term patterns also differ markedly in many cases. The reason that agents behave so differently is due to the extreme values of $NatShare$. The ratio of the two groups is an important parameter in a Schelling model (see Section 2.4) and in the case of adaptive agents, extreme ratios in which one group outnumbers the other by 8 or 9 to 1 cause the otherwise stable patterns to break down. When this situation occurs, natives and migrants differ in their behaviour, suggesting that the fact that migrants come in later on does change the short-term behaviour of agents. The control of $E = 0$ is roughly a mirror image with some variation likely due to the larger range of cases of $0.8 \leq NatShare \leq 1$, but the same overall pattern. When $E > 0$, native minorities can adjust their tolerance upwards whereas migrant minorities are slower to do so and may not reach the same levels at any stage. Natives are also more tolerant when they are the majority and are exposed to migrants at a slower rate. This seems intuitive, yet again cannot be confirmed due to the still-changing tolerance values at t_{max}. When natives are the large majority, they suffer a shock in tolerance (as they do when the model converges), but are slower to recover. This is due to the very slow rate of change of tolerance. When Δ_f is so low, agents' response rate effectively slows down. The relationship between levels of Δ_f and the time required to reach convergence is not linear. Whilst I cannot draw definite conclusions from non-convergent cases, the preceding graph has highlighted the sensitivities of the model with regards to extreme values of $NatShare$.

3.3.4 Cross-sectional analysis

In this section I analyse the dependent variables listed in Table 3.6 across a total of 15,000 independent simulation runs, drawing parameters from the distributions in Table 3.1.

Fig 3.6 shows a scatter-plot of the average final tolerance of each group against the native-share initial condition ($NatShare$), subdivided by treatment. The first column shows the tolerance of natives $\bar{f}_G$ and the second column depicts the tolerance of migrants $\bar{f}_B$. Outcomes from the five immigration treatments from Table 3.5 are each shown on a separate row.

Because both natives and migrants share the same decision-rule for adapting their tolerance (algorithm 3.2), my initial intuition was that all the graphs would simply be mirrored. The control treatment with no migration, $E = 0$, shows this mirroring pattern for both native and migrant agents: when natives are in a minority, their tolerance is more varied, dropping to 0.8. The same behaviour is observed for migrants when they are in the minority.

When migration occurs only once ($E = 1$), natives and migrants differ markedly in their tolerance behaviour. Part of the pattern is still mirrored: natives that constitute the vast majority of agents ($NatShare \geq 0.7$) are very tolerant, as are migrants when $NatShare \leq 0.3$. When natives are in majorities smaller than 0.7, the native tolerance splits: a large number of cases see very high native tolerance, and very low tolerance. Medium levels of $\bar{f}_G$ are observed throughout. A part of this pattern is reflected in the earlier time-series of native tolerance in Fig 3.4. Migrants don't diverge in their behaviour as much, although their tolerance starts to vary more as well, never dropping below 0.6. Migrants in minorities cope better than natives in minorities when $E = 1$.

This pattern does not apply for the cases $E > 1$. When $E = 4$, both natives and migrants are very tolerant with the exception of one outlier each, where tolerance is at or near zero. This means that the tendency to the extreme in these cases has

Figure 3.6: Scatter plots, by treatment, of native tolerance ($\bar{f}_G$) and migrant tolerance ($\bar{f}_B$) against native share of the population ($NatShare$), in steady-state at $t = t_{max}$.

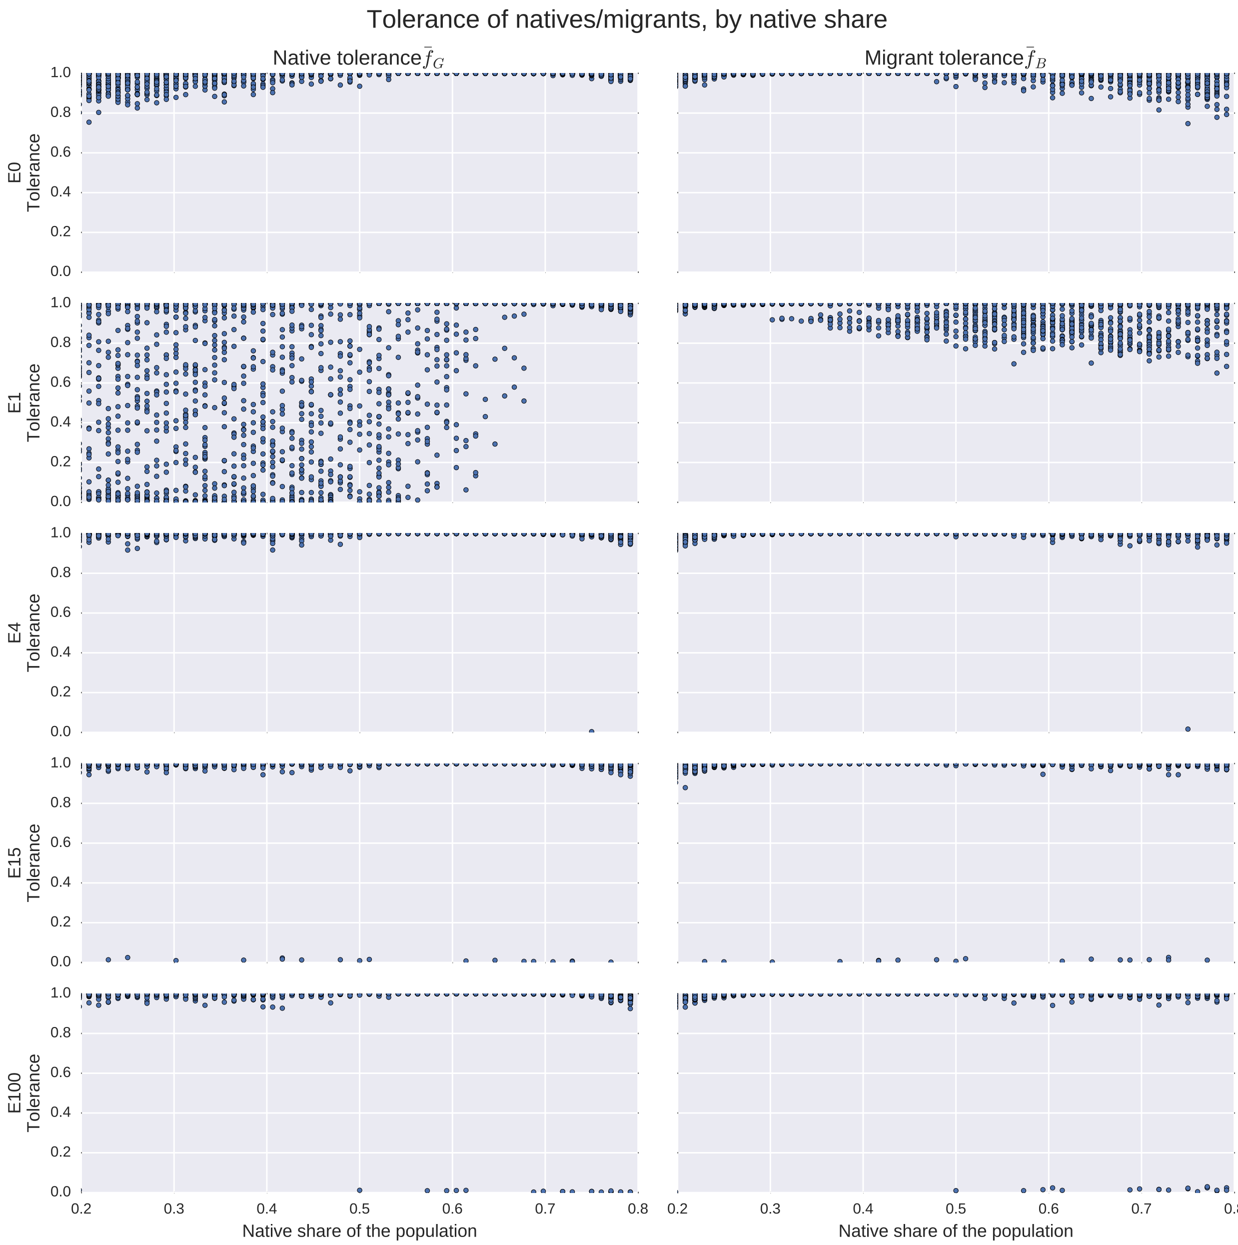

shifted to the other end. When $E = 15$, more cases of intolerant natives and migrants occur throughout the range of $NatShare$. When $E = 100$, those low-tolerance cases only appear at $NatShare \geq 0.5$. Thus, when the number of migration waves is very high and natives form the majority of the population, more cases of intolerance occur. Tolerance levels always verge on the extreme ends of the scale.

The different types of migration flows modelled by the treatments in Table 3.5

do not significantly affect the broad functional relationship between tolerance ($\bar{f}$) and native-share ($NatShare$) if $E > 1$. The high variance of native tolerance at $E = 1$ is likely due to the fact that the one-off migration wave is so disruptive that for large parameter ranges, natives do not recover their tolerance.

High tolerance within very small groups is due to the fact that with so few agents, no coherent group can form, and thus all free-moving agents become more and more tolerant as they have no homogeneous neighbourhood to escape into. This line of reasoning is intuitive for migrants in general, since they arrive in smaller batches. The high variance of tolerant minority migrants is down to their initial placement: depending on where their clusters are located, they may find a cluster large enough to ensure lower tolerance levels: intolerant migrants will stick to the cluster and if they are inside of the cluster, they will stop contact with outgroup members, thus reducing their tolerance. Migrants that have higher tolerance levels can roam the grid on their own: many neighbourhoods are satisfactory to them, and being in an unhappy contact situation, they will first try and move away before dropping their tolerance. Moving away presents them with a bigger range of potential satisfactory locations compared to intolerant agents. This means that tolerant agents will only by chance end up increasing the size of an intolerant cluster. It is more likely that they will settle in other areas of the map, simply because there is more of it.

I recall that agents are utility-satisficing, not -maximizing. That is, a tolerant agent that accepts a 20% outgroup-share in their neighbourhood will view a 20% share the same way as a neighbourhood with a 5% outgroup-share. Thus, the locations that would be picked by a maximizer, those around the intolerant clusters of ingroups, remain vacant and are picked only by chance. The agents that can sustain their intolerance reside at the centre of such ‘intolerance hot-spots’. These clusters must be large enough to sustain themselves. Too small, and the fringe agents will become more and more tolerant (their access to the cluster satisfies their tolerance requirements, but their outside access provides positive contact), and thus tolerate more and more outgroup members. As these move in, the intolerant hotspot disintegrates, starting outside and 'eating' inwards: the cluster becomes more and more tolerant until it breaks up. A few intolerant agents will move away and smaller intolerant clusters form, but they too break apart. The smaller the clusters, the quicker their disintegration.

3.3.5 Bimodality analysis

Both the cross-sectional and the time-series analysis indicate highly polarised tolerance levels. Fig 3.7 shows a histogram of tolerance values f_i across the population at the end of representative simulation runs. In these cases, tolerance values are concentrated at both extremes of the distribution, and the distribution is bimodal. Depending on the parameters such as Δ_f and $NatShare$, the split can vary between 30-70 and 70-30, with less than 10% of agents taking more moderate tolerance values. Thus parameters can determine the extent to which a population leans to the very tolerant or very intolerant, but the overall pattern remains that of a deeply divided society.

Figure 3.7: Histograms of the three most common distributions of final tolerance levels f_i. Intermediate tolerance is infrequent when tolerance levels are polarised.

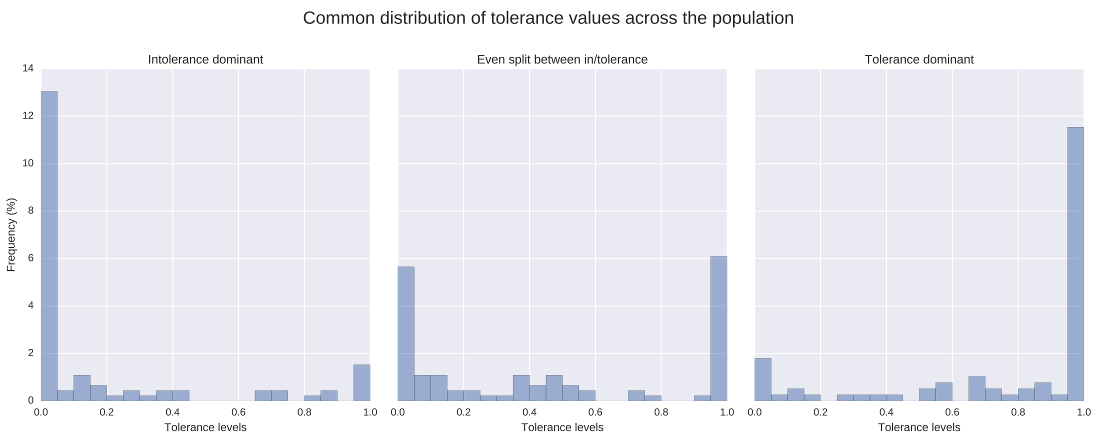

Fig 3.8 shows a scatterplot of the bimodality coefficient of tolerance B_f against the native share of the population across all simulation runs. Regardless of the experimental setup, the bimodality coefficient is almost always above the critical

value $B_f > \frac{5}{9}$, denoted by the red horizontal line. Due to their similarity, $E = 1$ and $E = 4$, as well as $E = 15$ and $E = 100$ were grouped together. Bimodality is lowest when native share is either $< \frac{1}{3}$ or $> \frac{2}{3}$ of the population. In cases with migration, bimodality drops earlier compared to the control. In a large number of cases, bimodality is nearly at 1 for the mid-range values of native share, illustrating its strong polarising effect on the population. The sharp drops in bimodality near the critical value are caused mainly by very low values of Δ_f. When the rate of change of tolerance is near-zero, it increases the chance that the tolerance values change too slowly to ever hit the extremes. That said, these lower values are still well above the critical threshold.

Figure 3.8: Scatterplot of the bimodality coefficient of the tolerance distribution β_f at the end of simulations, against native share of the population ($NatShare$). Treatments were merged here due to their similarity. The critical value $\beta_f > \frac{5}{9}$ is denoted by the horizontal red line.

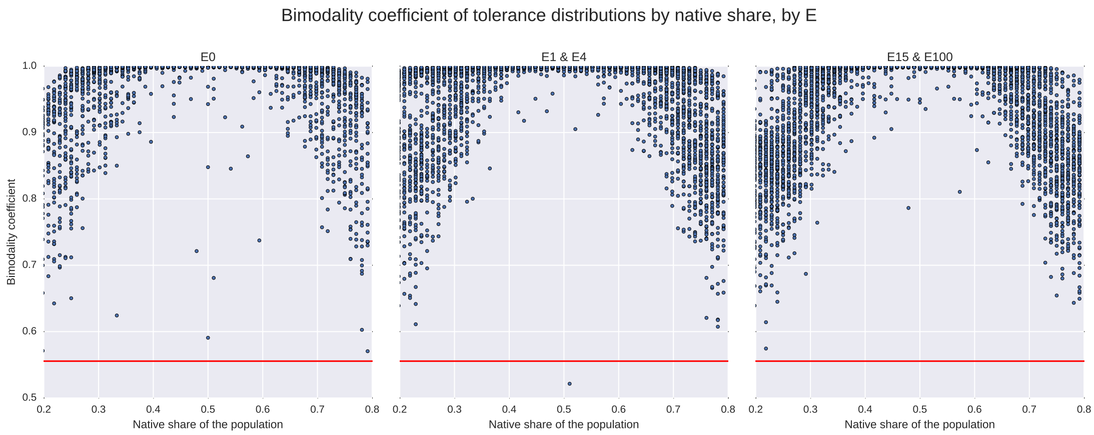

The bimodality in this model is a crucial indicator that the moderate levels of tolerance are not stable. Agents will tend towards either extreme, sooner or later- even when the value of Δ_f is very low, requiring up to forty consecutive turns to change tolerance so that one additional agent is tolerated (or not tolerated) in the neighbourhood. The trend towards extremes is driven by the self-enforcing nature of the neighbourhoods, as described in section 3.3.4 above: intolerant agents will segregate until they are happy, and start dropping in tolerance once they lose any contact to outgroup members. Tolerant agents will reaffirm their positivity by increasing their tolerance as a result of being happy whilst being in contact with the outgroup. Extreme tolerance is more stable than moderate tolerance. The cases in which the intolerant extreme dominates, segregation has established large enough clusters that sustain itself, and density levels have provided enough vacant buffer zones that remove the enforced contact situation in high-density settings. In these scenarios, tolerant agents reside on the fringes, but are not numerous enough to penetrate the intolerant core of the segregated cluster.

A divided society of agents is typical for a Schelling model, since agents are intrinsically homophilic by construction. However, divisions in a Schelling model are based on colour, or in this case, native or migrant status. The segregation of *tolerance* in this case is higher: Figure 3.9 shows the Moran's I of spatial autocorrelation for both colour-based segregation M^c (as typically measured in a Schelling model) and for tolerance-based segregation M^f. Again, the results are grouped into both ends of native share values and broken down by immigration treatment. The long-term segregation levels for out-groups (colour) are consistent throughout all cases, including the control (Figure 3.9a,f), never reaching $M^c = 0.2$. The values are not much higher than the values observed when movement is random (i.e. $0.1 \leq M^c \leq 0.15$). Schelling models have a baseline M^c value because some segregation exists by pure chance of agents' position at any point in time, suggesting that adaptive agents can circumvent segregation of colour to a large extent.

By contrast, segregation of tolerance attitudes, M^f is higher at the end of each simulation for most treatments of E and levels of $NatShare$. The exception is $E = 1$ and low $NatShare$, Fig 3.9 b which sees both colour and tolerance segregation at similar low levels. Segregation is much lower compared to traditional Schelling models because segregating is no longer the only way agents can mitigate unhappiness. Because moving away only occurs in once instance (contact + unhappiness) and ad-

Figure 3.9: Time series, by treatment, of the segregation levels of colour M_t^c (black), and tolerance M_t^f (yellow), filtered by extreme values of the native-share initial condition ($0.3 \leq NatShare \geq 0.7$). The error bars show the 95% confidence intervals of the mean of M_t across simulation runs. The dashed line at t_{mig} marks the end of migration waves.

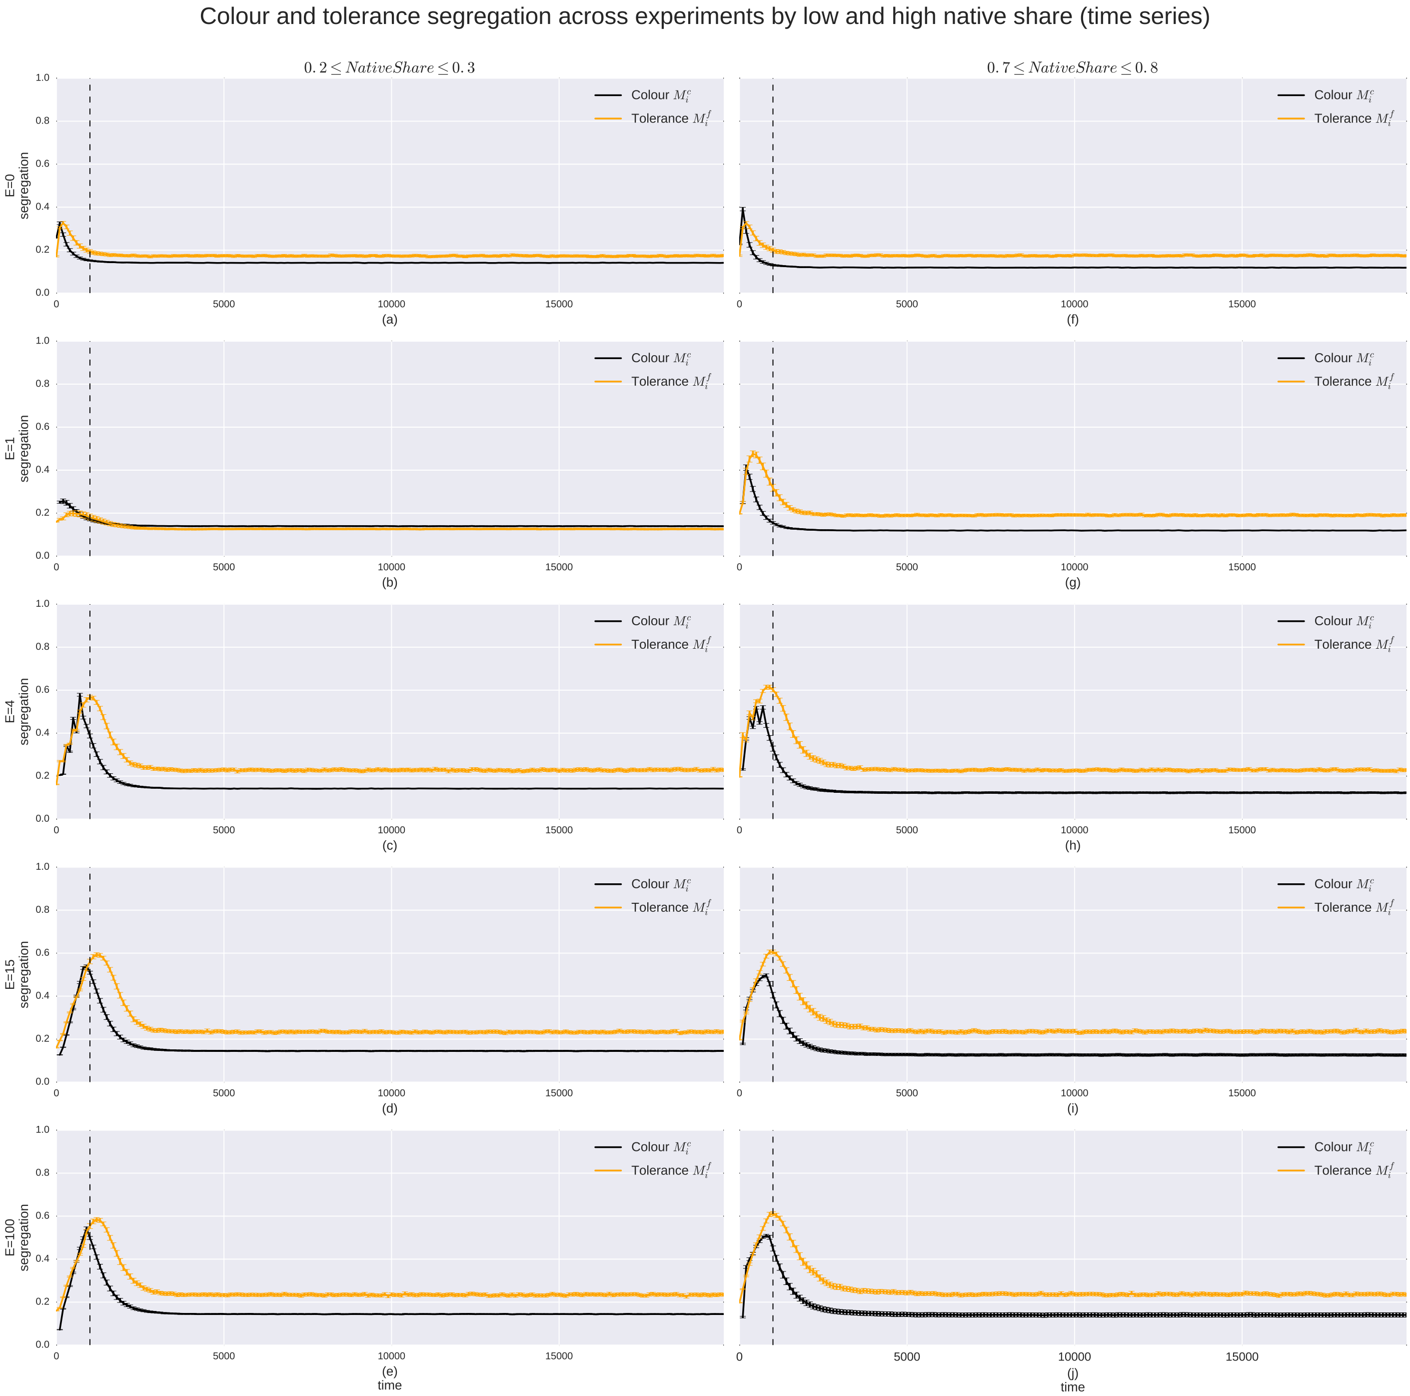

justing tolerance levels occurs in two instances, segregation is no longer the required proxy for tolerance of agents.

The differences between M^c and M^f are more pronounced when $NatShare$ is high (Fig 3.9 h-j). In the case of $E = 4$ (Fig 3.9 h), each migration wave is a visible shock to the existing tolerance level. Each drives M^f and M^c up. The same pattern is true of $E > 4$ (Fig 3.9 i,j), but the smaller waves leave less distinct marks. As seen previously with tolerance developments in Fig 3.4, the values change most during the arrival of migration waves, and settle once the final wave has arrived. Convergence of segregation values is reached before $t = 5000$, where cases of $E > 1$ and high $NatShare$ (Fig 3.9 h-j) need slightly longer than cases with low $NatShare$ (Fig 3.9 c-e). Long-term levels of tolerance segregation are at or above 0.2 when $E > 1$, whereas no migration shows lower levels of M^f.

The higher level of attitude-based segregation is interesting because agents do not actively seek out tolerant or intolerant neighbours; in fact, they are oblivious to the tolerance attitudes of their neighbours. The movement rule, as described in 3.1 and 3.2, drives adjustment of attitudes, but not an open choice of out-group neighbourhood, as is the case with agent colour. The connection between diverse neighbourhoods and adjustments of tolerance causes an unintended, much larger segregation than that of out-groups. Previous work including segregation of preferences have found the same effect (see Urselmans (2016) and Hatna and Benenson (2015b)), but the scale of M^f is much lower in this adaptive model (see Hatna and Benenson (2015b) for more insight into M^f in a non-adaptive model). Segregation levels of tolerance drop after the initial spike caused by migration waves (or initial settlement in the case of $E = 0$) because the tolerance levels adjust. As we know from the previous results, agents will tend towards the extremes of tolerance. This in turn means that the differences in tolerance between all agents become smaller, and thus the segregation measurement does not detect as large a difference. For instance, a group of agents that share a tolerance range of 0.7 and 0.9 will not register as a significantly divided group, since the difference is only 0.2. Because intolerant agents reside in segregated clusters, the mixing of tolerance groups is rare.

As discussed above, the tolerance adjustments cause a domino effect: tolerant agents will be surrounded by slightly less tolerant agents, who then turn tolerant as a result of positive contact. The reason why segregation levels spike during migration waves is two-fold: firstly, there is an artificial upheaval caused by the arrival of new agents that have a normally distributed tolerance range. They are yet to adjust their tolerance, and thus cause an uptick of moderate or low-tolerance values that are distinct from the established patterns. Secondly, for many agents migration waves become the first time of outgroup contact: movement increases and segregation increases with it, as intolerant agents will try and re-settle somewhere else. After the final migration wave arrives, tolerance levels start adjusting and as a result, segregation drops: both colour segregation, as agents become more tolerant, and tolerance segregation, because levels of high tolerance are more similar than levels of low and high tolerance.

3.4 Discussion

By implementing a version of the contact theory into a Schelling model of immigration, I was able to generate a highly divided society split between tolerance and intolerance towards out-group members. Contact theory may not predict a divided society, but its proposed mechanisms of positive contact versus negative contact as implemented here do lead to a society that is reminiscent of several Western societies that have experienced an increase in populist parties targeting immigration or immigrants as a problem to be tackled (see below). The model reveals several answers to the research question "how does immigration affect a host society and its migrant community", which will be discussed in turn below.

From the perspective of intergroup contact theories and threat theory, the results highlight the importance of minority and majority situations, which is corroborated by other studies (Hatna and Benenson, 2012). Migration affects the tolerance and segregation levels of a society, but minority-majority population shares mediates the effects strongly, especially in the short-term. The differences are particularly pronounced for more extreme values of the population share, and the effects can be both negative as well as positive. The importance of population share is in line with Allport (1954)'s suggestion that the respective standings of each out-group in society plays an important role in contact situations. My model suggests that the mere size difference leads to logistic situations that affect inter-group contact. This result also resonates with some of the findings in the empirical literature. For example, Pettigrew et al. (2007) find that majority groups experience a greater decline in prejudice as the result of contact compared to minority groups. The status of groups influences the potential perceived threat (Stephan and Renfro, 2002). Dominant groups might fear a loss of privilege, the subordinate groups might worry about oppression(Stephan and Renfro, 2002, p.195).

From the perspective of political science, I have shown that behavioural rules based on existing social theory can give rise to a polarisation in tolerance towards out-groups, and self-assortment of the population into tolerant and intolerant clusters. Neither behaviour is built into the model from the outset, but rather these are emergent behaviours that arise from a subtle interaction between adaptive tolerance and movement of the population. Crucially, divisions of native-migrant groups do not necessarily predict divisions in tolerance. Rather, intermediate tolerance levels are inherently unstable when movement and tolerance-adaptation interact. Sometimes individual agents transition from one extreme to the other, but most agents remain either very tolerant or intolerant. That medium levels of tolerance are an unstable spatial pattern was unexpected. My intuition was that the largest section of society is moderate, following a bell-shaped curve of tolerance distribution. However, findings from the most recent European Social Survey have found the same development of bimodality: across Europe, attitudes on immigration are becoming more and more polarised (Lambert et al., 2017).

This assortment of the population into communities of similar tolerance is highly reminiscent of the 2016/17 political landscape in the UK, the US and the many European countries that have experienced a surge in populist parties, although the trends towards negative views on immigration have been observed for years prior (Drinkwater et al., 2013).

These new developments have shifted the divide between economic left and right to more sociocultural divisions. In Britain, those who voted to leave the EU in the 2016 referendum were characterised by social conservatism, nationalism and low levels of political trust, whereas remain voters were more likely social liberals, cosmopolitan and high on trust values (Ford and Goodwin, 2017). Similar divisions are visible for Trump and Clinton voters in the US. Social and economic ideologies change how voters perceive social and economic issues (Crawford et al., 2017). Non-economic issues have become increasingly important for political parties in the West (Inglehart and Norris, 2017), and populist parties and candidates appeal on the basis of fears about immigration, sovereignty, and security. Both liberals and conservatives are subject to increased prejudice towards the respective political out-group (Crawford et al., 2017), but the rise in populist narrative builds on immigrant narratives especially in the Netherlands and in the Brexit referendum in Britain (Ford and Goodwin, 2017). The Leave majority was highest in areas that were the least diverse or featured high numbers of working-class voters; but also in areas which had experienced rapid demographic change as the result of immigration in the past ten years (Ford and Goodwin, 2017). These empirical findings speak to the strength of the contact theory.

Moreover, the *mechanism* that causes self-assortment along the tolerance axis also has a plausible real-world analogy. Within the model, polarisation of tolerance is caused by the shield-and-buffer dynamic described in Section 3.3.1 which prevents clusters of intolerant agents from one group from being exposed to out-group agents.

This shield-and-buffer dynamic might offer one potential explanation for the political bimodality found today. Cosmopolitan areas of a country are generally populated by people that are more tolerant of migrants; and migrants that reside in these areas are generally tolerant in turn (Goodwin and Heath, 2016). However, if the migrant diaspora exceeds a certain size, it could potentially sustain a sub-culture that is not dependent on integration with the host population. If non-migrants that live outside these areas have no direct contact with migrants, or the migrant fraction of the population is too small to provide sufficient contact situations, and peoples'out-group tolerance increases or decreases as described by contact and threat-theory, then polarisation would ensue, just as it does in my model.

Empirical studies seeking to support either contact or threat theory have been inconclusive (see Section 3.1), and more recent meta-studies lean slightly in favour of threat theory, although results depend heavily on operationalisation of variables and unit of analysis (Kaufmann and Goodwin, 2016). Implementing an Agent-based model of the contact theory has shown that the individual-level mechanisms suggested in the contact theory are sufficient to generate a macro-level pattern that reflects stylistic facts about migrant societies. The impact of immigration on a host society is mediated by the rate of change of migration, of tolerance, and the minority-majority relations between the two population groups.

4 Chapter 4

In this chapter, I move on from the physical segregation of (native) citizens and migrants and shift attention to social norms in an immigrant society. Norms are a subject of interest on their own, but the growing proportions of migrants in the West has highlighted their importance particularly when different norms clash. Norms do not just change what we do, but also change how we think about what we ought to do, and even what we *ought to think* (Brennan et al., 2013). Social norms have been the subject of interests of computer modellers (see Chapter 1) due to their inherent complexities: they are unwritten, widely adhered, can emerge and be sustained without enforcement (Willer et al., 2009), spread quickly and become obsolete (Mitchell, 2009).

Norms are very relevant to migration for a number of reasons. For instance, they they are important for the contextual neighbourhood effects that I have explored in Chapter 2. Neighbourhoods with positive social norms have higher satisfaction and less movement (Van Assche et al., 2018). Social norms can lead people to silence themselves in order to 'fit in' (Sunstein, 2003). This is of particular relevance for the study of public opinion- both in relation to migration, but also in general. If people silence themselves for long periods of time, perceived consensus can be misleading (Shamir and Shamir, 1997). The recent re-emergence of populism has raised questions as to what precipitated it (Groshek and Koc-Michalska, 2017). The internal pressure to conform offers an explanation for such perceived sudden changes of public support.

The focus of this chapter is on a particular kind of misperception of public support for a norm: pluralistic ignorance (PI). PI describes a situation in which a group of people erroneously believe in majority support for a given norm or attitude. For instance, college students might believe that most students are in favour of excessive alcohol consumption, and it is only themselves that do not agree with the practice- when in fact, a silent majority disapproves of the practice. Yet, in the collective belief that alcohol consumption has majority blessing, students do not signal their discontent (see Brennan et al. (2013) for more details on the college drinking studies). PI has been used to explain sudden shifts in public opinion (Shamir and Shamir, 1997), which are normally slow to change. Attitudes change slowly, yet public opinion can sometimes undergo rapid shifts, such as the growing support for same-sex marriage (Perryman et al., 2017). Pluralistic ignorance can give rise to widespread misperceptions about public opinion consensus, so that when attitudes shifts occur, they have in fact been occurring for a long time- but the perception of majority consensus has led people to misrepresent their personal views or led them to believe that others were truthful in their public attitude positioning (Shamir and Shamir, 1997).

I investigate the phenomenon of sudden shifts in public opinion or rather, support for popular norms. I conjecture that, as suggested by previous scholars (Kuran, 1995), (Shamir and Shamir, 1997), pluralistic ignorance is the reason why some shifts appear so sudden. Based on the existing P.I. theory, I suggest a mechanism why pluralistic ignorance can occur, and which circumstances can facilitate its occurrence and lead to vulnerabilities to the false consensus effect: the mistaken belief in a consensus which does not exist.

The goal is to determine what behaviours lead to a high level of PI, and what behaviours can mitigate PI occurrences. Agent behaviour is modelled on the basis of existing findings in the literature. In the context of the migration society modelled in the previous chapters, assume a society as bimodal as those generated in Chapter 3. In this version of the model, there are no migrants and natives. Instead, there are those that adhere to a norm of tolerance, and those that do not: an ideological segregation of norms. Crucially, agents can believe privately in one norm, but publicly express their support for another norm. The model is again based on the Schelling framework to simulate the social environment.

Exposure to false consensus effects can explain sudden shifts such as experienced in the case of same-sex marriage (Perryman et al., 2017), but it also opens up possibilities for populists to reach out to those people that silence themselves because they don't want to appear as if they violate existing norms of tolerance or support for multiculturalism (Blinder et al., 2013). Populism is not a new phenomenon and has experienced waves of support before (Norris and Inglehart, 2018). A core defining feature of populism is that it challenges established institutions and/or authorities, and presents the status quo as something that is no longer feasible (Norris and Inglehart, 2018). Attacking the legitimacy of the prevailing system and set of beliefs is a way to reach people who silence themselves. Finding the prerequisites for such fertile ground for populists is possible by studying when pluralistic ignorance can arise.

4.1 Introduction

Typically, social norms are modelled as spreading or emerging usually by means of contagion, imitation, learning, and coercion or through rational-choice and expected utility-maximising behaviour (Beheshti and Sukthankar (2014), Centola et al. (2005)). Wang et al. (2013) introduce an opinion-dynamics model adapted to incorporate pluralistic ignorance. In their model, a single agent rejecting the status of pluralistic ignorance can lead to a complete change in opinions of the entire group, if the agent's opinion is firm and its neighbours are unsure about their own attitude (Wang et al., 2013, p.247). Nevertheless, the occurrence of pluralistic ignorance presents a puzzle for theories of social norms, because it is unclear how it occurs in the first place (how it differs from individual psychological error) and how it can spread and persist.

Agent-based modelling enables us to look at the emergence of pluralistic ignorance from a bottom-up perspective. By giving the individuals rules of behaviour on a micro level, the model can give rise to macro-level patterns of behaviour that can illustrate how emergence can occur, i.e. which preconditions are needed and what rules can recreate the stylistic facts that have been established in the literature. As Centola et al. (2005) note, pluralistic ignorance theory has several shortcomings, one of which will be addressed in this chapter. The existing theory makes no explicit assumptions about the group element of pluralistic ignorance, an important part that I will demonstrate to be crucial to define and understand pluralistic ignorance. I show how PI can arise as a result of a desire to be with others with similar public behaviour, without the need for switching norms or sophisticated information exchange mechanisms. Secondly, I show that the restriction of movement through social space can explain how pluralistic ignorance can persist even when many people are truthful in their public expression of their private views.

The aim of the chapter is two-fold. Firstly, it aims at contributing to the existing theory of pluralistic ignorance by examining the crucial *group-conditions* that are necessary for pluralistic ignorance to unfold. Secondly the chapter demonstrates how pluralistic ignorance can emerge from very few and simple rules that can lead to complex macro-patterns of norm adherence and rejection, the rule being that people desire to be with others with similar public behaviour- there is no need for attitude-switching or information-exchange mechanics to explain the emergence of pluralistic ignorance. The relevance to immigration societies is discussed at the end of this chapter, Section 4.6.

The study of pluralistic ignorance dates back to Katz and Allport (1931) who are widely credited with first developing the notion of pluralistic ignorance. Since the first definition of PI by Katz and Allport (1931), the concept has been expanded and has led to many different nuanced interpretations of PI. Brennan et al. (2013) cite Allport saying that PI is "[...] a situation where a majority of group members privately reject a norm, but assume (incorrectly) that most others accept it" (Allport, 1954). For example, a norm could be table etiquette: the use of the fork with the left hand. A scenario of pluralistic ignorance would be: a family gathered around the table; everyone using the left hand to operate their forks. However, every member of the family secretly wishes to switch hands. But they all believe that they are the only ones wishing to deviate from the (perceived) norm. The unwritten nature of norms makes these situations more likely. It is reasonable to assume that neither parent at the table has a written version of the etiquette. Because people use heuristics to guide them through social comparison (Kahneman, 2011), assumptions are not always the result of informed decisions.

In the realm of Political Science, pluralistic ignorance received intermittent attention in studies of public opinion. Public opinion can be the result of widespread pluralistic ignorance, explaining sudden shifts in public opinion (Lerman et al., 2016). Survey respondents for public opinion polls are aware of norms surrounding what they are supposed to be or think, and will feel pressured to answer accordingly (social acquiescence bias (Krosnick, 1999)). This bias can be the result of the same process that creates pluralistic ignorance: individuals wrongly assume that a given viewpoint is in fact supported by 'the majority', and will give biased answers in a survey.

Voters are subjected to norms such as the norm on whether one should vote or not, or which way one should vote. Generally, the norm in most OECD countries is that a citizen should vote, but some subgroups disregard the country-wide norm and replace it with a different one that discourages voting (Labovitz and Hagedorn, 1973). The study of public opinion and attitudes relies to a great extent on survey responses about policy preferences and the salience of issues (Stimson, 2004), both of which are subject to a lack of information and little willingness on the part of people to obtain that information, thus relying on less accurate but easier methods of interpreting signals (Kahneman, 2011). Voters will make decisions in the absence of much accurate information, and will resort to easier methods to reach a state in which they feel qualified to make a choice. The situation bears resemblance to the social comparison process of the fork-wielding family: having incomplete information, voters rely on inferring the missing bits by observing others. Absence of visible defection is assumed to signal compliance. For the public sphere as opposed to the situation at the family dining table, the role of the media is relevant as well. It can provide access to information that is otherwise not obtainable by individuals. However, when media provides a narrative for a given topic, people tend to disregard their personal experiences (Andersson et al., 2017). With an increased reliance on indirect sources via traditional media and social media, people are vulnerable to misreading cues- these can be intentionally misleading, or interpreted wrongly.

Pluralistic ignorance is a group phenomenon in that it is analytically distinct from the psychological process of deriving attitudes from behaviour (O'Gorman, 1986). But without the individual processes of misreading social cues, PI cannot exist. This makes it a very suitable subject for a bottom-up agent-based approach. By modelling individual behaviour of which we have empirical proof, we can study the emergence and persistence of pluralistic ignorance resulting from the collective of actions of individuals. Before laying out the agent-based model of pluralistic ignorance, I shall define the minimum requirements for PI in response to the existing debate around pluralistic ignorance and what it constitutes. Definitions of the phenomenon vary widely and capturing it in a computer model requires greater specificity.

4.2 Defining pluralistic ignorance

There is no single agreed-on definition of pluralistic ignorance in the literature (Centola et al., 2005). One of the intuitive examples often cited to explain pluralistic ignorance is Hans Christian Anderson's tale of the Emperor's new clothes. The respected and acclaimed emperor is fooled by a group of rogues into believing that their (in fact) non-existent new robe for the emperor is real. The rogues persuade the emperor that stupid people can't see the robe- and thus the emperor, afraid to admit that he can't see the robe, pretends that he can see it and that the robe exists. The citizens fear the emperor and pretend to be amazed by the garments- until an innocent child comes along and laughs at the naked emperor, and the spell of pluralistic ignorance is broken. The citizens realise that they had all been wrongly assuming they were the only one to see the naked emperor and join in with the laughter. The fable is appealing in that it is very easy to imagine. As strong as the intuition is, it is also vague and does not include all conditional attributes for pluralistic ignorance.

Because the Emperor's New clothes example is so ubiquitous in the literature of pluralistic ignorance and norms, I will use it to put forward the theoretical argument that different majorities need to be considered when trying to explain pluralistic ignorance as a group phenomenon. The fractionalised nature of pluralistic ignorance literature has led to the use of different examples and case studies that are not always comparable. I will thus use the Emperor's tale to ensure that the context of PI is clearly defined before applying it to real life cases.

Consider the following situation. The emperor rules over a court of 30 people who are all equal in their status, influence and visibility. All members can see each other and directly observe behaviour. Now imagine that five people pretend the emperor's clothes exist and these people also assume that everyone else believes it. The remaining 25 do not share that social reality: they assume everyone does *not* believe the clothes exist, and act accordingly: they laugh. It seems counter-intuitive to classify that as pluralistic ignorance, even though those five people form a group of PI. What if not just five, but 15 people believed in the pretence? Imagine that for some reason, the other 15 people sincerely believe the clothes exist. All 30 people signal that the clothes exist, and half of them are sincere in doing so. The assumed consensus of existing clothes (based on visual confirmation of signals) is actually not a majority; it's 50% of the population. Yet 100% of the population makes the (erroneous) assumption of a majority. There are three questions related to the group of PI that arise from this tale:

1. What is the minimum group size for pluralistic ignorance?

    1.1 What are the conditions of behaviour within the subgroup?

2. Is a social group, in which its members *cannot* directly observe one another, a valid group for pluralistic ignorance?

Deciding what the minimum group size for PI should be affects the range of cases that it can be applied to. For the study of public opinion, group sizes are naturally large. For sociological experiments, group sizes can be very small. Which size should it thus be?

A source of confusion is that in a PI scenario, there are two different kinds of majorities. The first majority is that of the assumed popular opinion ("the emperor's clothes are real!"). The second majority is the fraction of people (in this case, at least 16) within a defined group (the emperor's court) that has to believe in the majority opinion, and they have to be wrong. Not all definitions of pluralistic ignorance make that distinction.

The most basic condition for pluralistic ignorance is the requirement of a shared cognitive error by a group of people. PI definitions can be categorized into either focussing on the collective, the group and the environment in which pluralistic ignorance occurs, and those focussing on the individual, trying to explain the psychological reasons for individual failings. Allport's definition (quoted earlier) implies that a precondition for pluralistic ignorance is that a majority of a group has to suffer from the cognitive error, and that the assumed norm consensus must be shared by 'most' people. This detail is not present in many other definitions.

Definitions focussing on the individual stress that people have to have a distinct private attitude and public display of behaviour, and that it often involves wrong assumption about others. This interpretation is often attributed to Miller and McFarland (1991), stating "[p]luralistic ignorance is a psychological state characterized by the belief that one's private attitudes and judgements are different from those of others, even though one's public behaviour is identical [...]"(p.287). This definition describes the process of erroneous judgement that is required for pluralistic ignorance to occur. It differs markedly from the Allport definition, which specifies majority requirements of PI groups.

4.2.1 Group size and subgroups

A more specific definition including the size of groups was presented by O'Gorman (1986) stating that "[p]luralistic ignorance refers to erroneous cognitive beliefs shared by two or more individuals about the ideas, feelings, and action of others." According to O'Gorman, any pair of individuals (sharing their cognitive error) within any group would constitute pluralistic ignorance. This is not compatible with Brennan & Goodin's notion that a majority of a group has to share the cognitive error, unless a pair of people constitutes a majority in a group of three (Brennan et al., 2013). The notion of group size varies between definitions and research areas. From a public opinion perspective, the group can encompass large elements of society, or society as a whole (see for example Moy and Rinke (2012); O'Gorman and Garry (1976); and Kuran (1995)). Other studies mention smaller groups, such as social groups at university (Prentice and Miller, 1993), inmates in prisons (O'Gorman, 1986) or a religious community in a localized space (Schanck, 1932). Bjerring et al. (2014) offer an in-depth discussion on the variety of pluralistic ignorance definitions and conclude that many definitions offered in fact lead to overestimation of pluralistic ignorance as the definitions leave too much theoretical leeway. They offer the following alternative definition of pluralistic ignorance:

> [...] "Pluralistic ignorance" refers to a situation where the individual members of a group
>
> (i) all privately believe some proposition P;
>
> (ii) all believe that everyone else believes ¬P;
>
> (iii) all act contrary to their private belief that P (i.e. act as if they believe ¬P); and where
>
> (iv) all take the actions of the others as strong evidence for their private beliefs about P.

This definition contains both individual-based (iv) and group-based elements (i-iii). However, I argue that the notion of "all" individuals and "all" actions is not justified for cases (iii) and (iv). It is not necessary for *all* members of a group to act contrary to their beliefs, nor that everyone, as stated in (ii), believes that anyway. If truly every single member of a relevant group has to follow the same pattern, these strict conditions are unlikely to ever be met in the real world. It is reasonable to assume that most authors will acknowledge that the absoluteness of all members is in fact more flexible once applied to the real world.

The importance of absoluteness instead applies to the individual cognitive process that can lead to pluralistic ignorance: people tend to think of 'everyone else' when they gauge public positions on an issue (Gunther et al., 2008). In other words, members of a PI group may believe that *everyone else* believes, but this is the actor's subjective view of the situation, not the objective assessment. While the individual views are part of the mechanisms of social behaviour that can lead to PI, they do not form part of the definition of the group of PI. Another way of distinguishing the two could be the individual psychological explanation versus the collective sociological explanation. The minimum group size should not be equal to the maximum group size. I believe that the notion of two or more people put forward by O'Gorman (1986) is the most useful minimum working example. There exists an ongoing debate among social psychologists whether dyads (pairs of two people) constitute a group or not (Williams, 2010). In the case of pluralistic ignorance, a minimum group size of three people (triads) is more applicable. Pluralistic ignorance stems from social comparison processes that are employed in order to reduce the complexity of a social situation. This complexity is a defining feature of triads (Moreland, 2010) and larger groups. Dyads are characterised by more intimate, intense and direct types of contact, a social context in which pluralistic ignorance is not normally defined in. Triads feature majority-minority relations, which dyads by definition cannot (Moreland, 2010). Pluralistic ignorance relies heavily on perceptions of majorities, so the minimum group size should be a triad of three individuals.

Thus, the first requirement for minimum pluralistic ignorance conditions is:

(a) Define a group as a set of individuals of a minimum of three members.

For pluralistic ignorance among the emperor's court, the minimum group size is three people. Three out of thirty people privately believe P: the emperor's clothes do not exist. Those three people also believe that everyone but themselves (29 others) believes $\neg$P: the emperor's clothes do exist. This private-public attitude inconsistency is widely accepted in studies of PI: the psychological process of individuals can entail the belief about *all* group members. Crucial is however the resulting behaviour: The three people will act as if they believe $\neg$P: they act as if the clothes existed. This addresses the second part of a minimal PI situation:

(a) Define a group as a set of individuals of a minimum of three members.

(b) Require that at all members of the PI group have the same public behaviour.

    Require at least a plurality of the members of the group to have the same inconsistency between their public behaviour and private beliefs.

These requirements translate into a political science related issue in the following way: the minimum group size ensures that virtually any political group can be studied, from country-wide or international public opinion and attitudes to the study of small groups of elites or various interest groups within a cabinet of a government, comprising of a handful of people. In the study of public opinion, display of public behaviour is translated into either the prevailing media narrative, opinion surveys or more recently, online sources such as social media engagement and twitter feeds. ‘Behaviour’ in the realm of public opinion refers to the display of messages, and crucially the absence of open dissent against a prevailing message.

The third question relates to the need to have physical or otherwise visible social interaction with others as a prerequisite for pluralistic ignorance. If yes, then groups will naturally be relatively small.

4.2.2 Mutual observability

Bjerring et al. (2014) argue that situations characterized by a “lack of observational interaction among agents in the relevant social group” (2014, p. 12) do not count as pluralistic ignorance. If there is no direct observation of others possible from which assumptions may be drawn, pluralistic ignorance is not present- supporting this claim is the wealth of previously mentioned examples akin to the Emperor’s New Clothes, all of which feature direct observation. When social interaction is not possible, other types of information gathering can take its place.

In any scenario with perfect information, pluralistic ignorance can, by definition, not occur. Any discrepancy of public behaviour versus private attitudes would be the result of an informed choice, not an uncertain assumption. Norm persistence is not dependent on what members of the group do but rather what they presume other group members do. Norms are sustained so long as people reveal their private attitudes only to a small number of people (Kitts, 2003), but pluralistic ignorance can persist even when there is widespread knowledge that most people in fact, misrepresent their private attitudes (Brennan et al., 2013).

Even if one concedes the strict precondition of direct observability, it raises the question of whether one must observe absolutely everyone else in a relevant group, or just a fraction: if social interaction is a prerequisite, pluralistic ignorance cannot occur in groups larger than the maximum number of people that can be observed simultaneously at any point in time. Widening the time-span, it would include the sum of all people that were observed to do something relevant to the private attribute in question. Images of others in day-to-day life are utilized using different sources of knowledge such as memory of observation or media consumption (O'Gorman, 1986). There is no convincing argument as to why knowledge of known behaviour must only be obtained through direct observation. O'Gorman (1986) posits that the "visible social milieu" of individuals and the "more distant and less visible social world of which that milieu is part" is both relevant to pluralistic ignorance, which is perhaps an analytically less pure but more realistic assumption. Pluralistic ignorance can result from 'small-world' networks: people have a limited social mileu, but will project their knowledge of that social circle onto the wider, indirectly linked or not linked nodes in the network (Lerman et al., 2016). Therefore, I reject the assumption of direct social observability. Instead, social interaction information can be obtained directly as well as indirectly. This can also incorporate the erroneous assumption that not observing defiance is interpreted as observing compliance with a norm (Brennan et al., 2013), or that public actions always reflect private attitudes of others.

Consider the bystander scenario, illustrated by Miller and McFarland (1991): an accident, such as a car crash, occurs and people arrive at the scene. No one has sufficient information to determine whether or not the scene is an emergency and warrants further help on their behalf. It is embarrassing to overreact and thus safer to err on the side of caution and show composure, even if that leads to a bystander situation which is ultimately not desired by any participant.

Whether or not a person truly believes that the situation is not an emergency, is ultimately not relevant: imagine an accident, and the first person to arrive at the scene is convinced that there is no emergency. The next person to arrive will use social comparison and investigate what the first person is doing, and then conclude that there is no emergency because that first person is not helping, and not signalling any distress. But the same result would occur if the first bystander was not a true believer: the first bystander knows of the norm of composure and will act accordingly. Pluralistic ignorance is not dependent on dispersal of information through observation. For example, people's estimation of public opinion towards an issue is influenced by their media perception (Perryman et al., 2017). Whether media coverage includes active acts of defiance against an existing norm (such as a public protest) can influence the estimate of majority public opinion of individuals.

To conclude, I define the minimum conditions for a group-focused definition of pluralistic ignorance as follows:

(a) Define a group as a set of individuals of a minimum of three members.

The set of individuals are either in direct or indirect contact through a (network) path

(b) Require that at all members of the PI group have the same public behaviour.

Require at least a plurality of the members of the group to have the same inconsistency between their public behaviour and private beliefs.

Indirect contact with different groups is a common source for information about social identities and can be substantial in the absence of any direct contact (Pettigrew et al., 2007). Inconsistency between public behaviour and private beliefs are well-documented in the context of eliciting truthful survey responses from people (Krosnick, 1999).

4.3 Scenarios of pluralistic ignorance in the model

The discussion on the preconditions of PI is of direct relevance not just for the theoretical concepts of PI, but also for the agent based model that I present in order to test the conditions for its emergence. In order to investigate whether one condition led to more pluralistic ignorance than another, I need to define what a group exhibiting pluralistic ignorance is characterized by in the first place. The definition that I offer might help to serve as a starting point towards a complete formal definition that includes the group conditions.

I am examining pluralistic ignorance under the strictest possible definition, requiring universal inconsistency, not just majority inconsistency. If I can show that desire to belong and movement leads to pluralistic ignorance merging with such a definition, my theory works under the toughest conditions. The rules are stricter than the theory in order to model the lower bound of pluralistic ignorance. The differences are described below. By simulating the harshest conditions for PI to occur, I prevent overstating its propensity to emerge and generating results with high levels of PI. Figure 4.1 visualises minimum PI groups and how pluralistic ignorance is captured in the model. Dark green tiles are pluralistic ignorance groups. The first letter is the private norm; the second letter is the public norm. An "AB" agent is thus privately following norm A, and publicly displaying support for norm B. Each scenario meets the minimum requirement for PI, but in every case, the group size of pluralistic ignorance is three.

The first scenario (a) contains three AB agents that are in a PI group. They all publicly adhere to B, yet all privately believe A. Three agents is the minimum amount for a group, and thus this scenario captures a PI situation. Scenario (b)

Figure 4.1: Scenarios of pluralistic ignorance

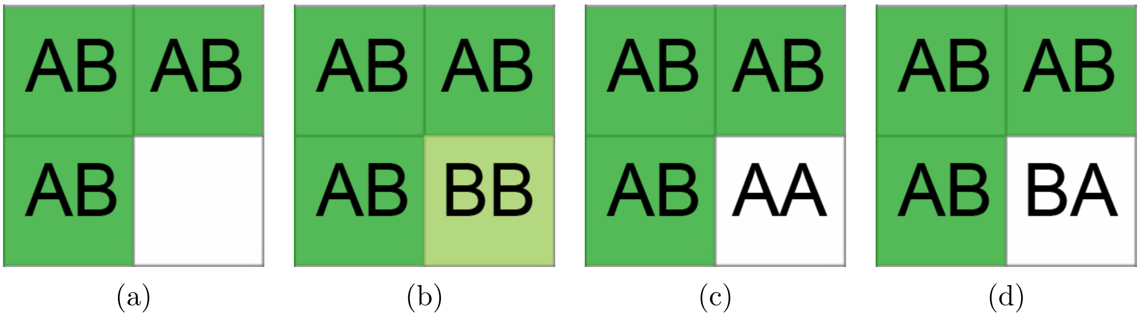

includes a fourth agent of type BB. This agent believes both privately and publicly in norm B. This agent is norm-consistent, or *truthful*. The algorithm employed to find PI groups will exclude this agent, because it does not share the same norm inconsistency as the other AB agents. As an individual, this agent can't be part of the majority of agents that form a PI group. But since the agent saw neighbours signalling publicly the same norm that the BB agent adheres to privately, the agent has moved into the group. As discussed in Section 4.2, the BB agent would in theory be part of the group that harbours pluralistic ignorance. Consider it as public opinion is believed to favour B, because agents have signalled so. But some agents will genuinely believe in B. In this scenario, 100% of agents publicly believe in B, and 25% do so genuinely. The majority consensus is an illusion to which the BB agent has indirectly contributed. To denote this, the tile is coloured light green.

Scenario (c) shows a case with an additional AA agent instead. This time, the AA agent does not contribute to the majority illusion of public adherence to B, because the agent signals A. But the three PI agents who all believe AB will have taken this agent into account: since they all privately adhere to A, they seek the company of an agent who publicly adheres to A. Consider public opinion to be believed to favour B by 75%, but 0% are sincere in this display. The majority is an illusion, but the AA agent has not contributed to it. This is also the case in scenario (d), where the fourth agent is of the type BA and does not signal norm adherence to the perceived majority. In this scenario, the AB agents and the BA agent attract each other. Each believes privately what the other publicly displays. This too does not constitute a part of the majority illusion, neither under the strict conditions of the model nor in theory.

In addition, the definition of a group in the model requires geographic proximity (adjacency on the grid), but the group must not necessarily be a rectangle-shaped cluster. As long as agents are adjacent to at least one other agent of the same norm inconsistency, they can form a group of pluralistic ignorance. For instance, a line of agents would constitute a group, because each agent is connected to at least one other agent: because direct visibility is not a requirement, agents can be part of a group through indirect network connections. Networks can give rise to and perpetuate pluralistic ignorance (Lerman et al., 2016).

Pluralistic ignorance captures the self-sustaining system of people feeding off other people's behaviour whilst signalling certain behaviour themselves. No one is a neutral observer to the system but always a participant, willing or not. The absolute notions of 'all people in a group' were discarded as too strict to be applicable to real life. A single helping bystander might not convince others that they should help if many others just stand and watch; not everyone at the court must pretend the Emperor's clothes exist, as long as enough others do. Individuals tend to assume that there exists their attitude and then that of 'all others', people don't tend to differentiate further (Gunther et al., 2008).

How does PI emerge and what behaviours foster its perseverance? The following section describes the implication of the definition decisions on the model elements, and how PI is captured empirically.

4.4 Method

The aim of the model is to simulate the emergence (and persistence) of pluralistic ignorance, testing several conditions for their propensity to generate pluralistic ignorance or stifle it.

A virtue of the model is that it contains so few parameters that are sufficient to give rise to pluralistic ignorance. Alternative (agent-based) models of pluralistic ignorance usually feature notions of "true-believers" and varying degrees of pressure exerted by different kinds of agents (see for example Centola et al. (2005)). I show that this is not necessary for pluralistic ignorance to arise.

4.4.1 The Model

The set of agents $A = \{a_1, \ldots, a_n\}$ are positioned on a toroidal 2-dimensional grid of size $S = 50 \times 50 = 2,500$ tiles W at time $t \in \mathbb{Z}$. Agents are assigned random tiles which they will occupy at the start of each simulation. Agents cannot enter or exit the grid. A number of agents is initialised at the start of each simulation which remains constant throughout. Each agent a_i has a public norm Ψ and a private norm Ω, both of which can have either value A ($\Psi_i = A$) or B ($\Omega_i = B$). Figure 4.2 shows a visualisation of the model. Agents have different colour visualisations depending on the norm that they follow. Public norm $\Psi = A$ is amber, $\Psi = B$ is dark blue. For instance, private norm $\Omega = A$ is yellow, public norm $\Omega = B$ is light blue. As a notational convention, when agents' norms are described, the private norm is listed first. Thus, an agent with norms BA has the private norm B and public norm A. Agents that hold differing private and public views are *norm-inconsistent*. AA and BB agents are *norm-consistent*.

Agents know both their private and their public norm, but can only see the public norms of other agents in their local neighbourhood. Private norms can never be disclosed or signalled to other agents.

Table 4.1: Constants

Constant	Description
$t_{max} = 10,000$	Maximum number of ticks per simulation
$S = 50 \times 50$	Size of grid

Figure 4.2: Example of a typical simulation run. Public norms are visualised on the top, private norms on the bottom. The private view shows distinct clusters of norms, the public view shows a more mixed population. States of the simulation at different times. Density 62%

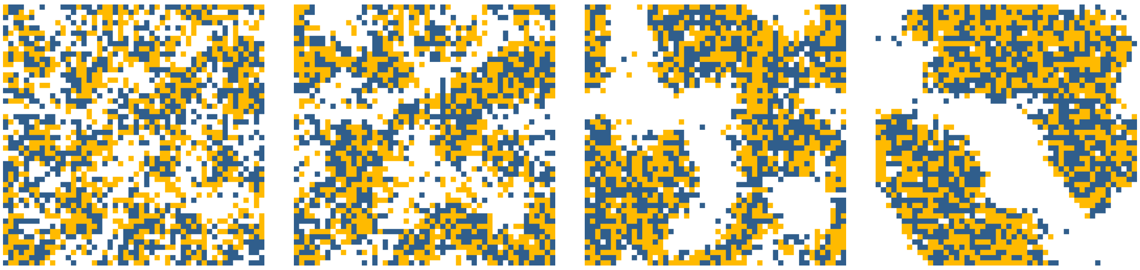

(a) Public Norms Ψ of agents at $t = 1, 10, 100, 500$, intelligent movement, $nbr = 3$

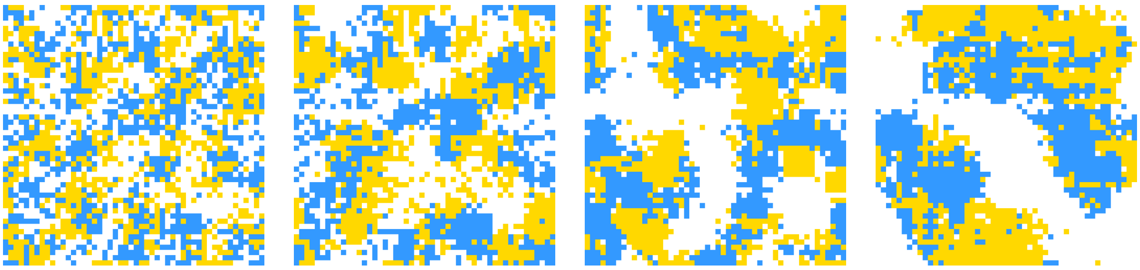

(b) Private Norms Ω of agents corresponding to the states above

Agents have a preference to reside in neighbourhoods that have the highest share (within that neighbourhood) of the agent's private norm. Since agents can only see public norms of their neighbours, they compare their own *private* norm to the public norms of their neighbours to determine the utility of each tile. Higher utility is always preferred. The agent's preferred norm followers s_i should outnumber the agents publicly following the other norm, d_i. The neighbourhood of an agent is represented by $N_t(a_i)$. Agents count themselves as 'sames', thus giving a slight advantage to 'same' agents. This is due to the propensity of people to assume that everyone else shares their beliefs; they would not see themselves as neutral bystanders.

$$
\begin{aligned}
s_i &= |\{a_j \in N_t(a_i) : \Psi_j = \Omega_i\}| + 1 && (4.1)\\
d_i &= |\{a_j \in N_t(a_i) : \Psi_j \neq \Omega_i\}| && (4.2)\\
& && (4.3)
\end{aligned}
$$

The differences of s and d is denoted by k.

$$k_i = s - d \tag{4.4}$$

Because agents will always prefer higher numbers of k_i (see equ. 4.4), they are utility maximizers, not satisficers. Let $N_t(a_i)$ denote the neighbourhood of an agent a_i at time t. Agents will count themselves when comparing the numbers of s_i to d_i in a neighbourhood. The size of the neighbourhood varies according to treatments. Those will be described in Section 4.4.3 (see Figure 4.3).

4.4.2 Movement treatments

I vary the movement and satisfaction behaviour of agents. Table 4.2 shows the different combinations of reasoning that agents can have. For each treatment, all agents on the grid will share this reasoning pattern; there are no differences between agents' movement behaviours for the duration of a simulation. For readability purposes, each movement rule has a name designed to reflect the reasoning patterns of agents. The terminology of 'intelligent' reflects that agents are able to adjust their behaviour based on their reasoning, and that their behaviour is not merely determined through random choice. As agents can conceal their private norm and display a different public norm, the reasoning patterns of intelligent, reasoning agents are based on sometimes faulty information. This may lead to suboptimal choices, but intelligently reasoned choices nonetheless.

In order to ascertain the effect of intelligent movement, I vary the agents' movement rules across four different experimental treatments summarised below in Table 4.2. These are described in detail below.

Table 4.2: The movement treatments

Treatment	Movement rule	Should I move reasoning	Where to move reasoning
Control	Random	Random	Random
Treatment 1	Intelligent	Intelligent	Intelligent
Treatment 2	Semi-Intelligent	Intelligent	Random
Treatment 3	Risky	Random	Intelligent

The control treatment induces random movement. The decision whether to move is assigned $p = .5$. If the agent decides to move, it will do so by picking a random vacant tile in the agent neighbourhood N_i, given by algorithm 4.1.

Algorithm 4.1 Move randomly

$L \leftarrow l \in \{N(p_i) : p_{i,t} = |a \in A : a_j = l| = 0\}$ ▷ all empty neighbourhood tiles
$l\star \leftarrow \textsc{ChooseOneAtRandom}(L)$
$p_{i,t+1} = l\star$ ▷ update location

Dissatisfied agents will try to move to an empty area within the agent's neighbourhood N_i. Agents cannot move beyond any tile within their current neighbourhood. This feature is meant to represent the social space (as opposed to geographical space) that the agents reside in. Friendship groups or networks such as work and family are durable and movement is costly: giving up friendships and finding new friends is a longer-term endeavour. The decision whether or not to move is given by algorithm 4.2.

These agents, labelled 'random', do not consider private or public norms in any way. Agents labelled 'risky' will also decide whether to move or not with a probability of $p = \frac{1}{2}$. Risky agents will not decide their future location randomly, but will do so by considering all vacant tiles in the neighbourhood that satisfy a norm plurality $s_i > d_i$, given by algorithm 4.3. The larger the plurality (or majority) k_i of norms in the neighbourhood $N(l)$, the more attractive is the location l. That is, agents are utility maximizers. Satisfied agents will not consider moving even if

Algorithm 4.2 Check need to move

```
NeedToMove ← False                                        ▷ default decision is to not move
for all l ∈ N(p_i) : p_i = |a ∈ A : a_j = p_i| = 0 do      ▷ check empty neighbourhood tiles
    s ← |{a_j ∈ N_t(l) : p_i = p_j}| + 1                  ▷ number of nearby public norm equals
    d ← |{a_j ∈ N_t(l) : p_i ≠ p_j}|                      ▷ number of nearby public norm differents
    if s > d then
        NeedToMove ← true                                 ▷ See 4.3 and 4.1
    else if d > s then
        NeedToMove ← false
    else p ← drawrandomlyfrom(0, 1)                       ▷ Probability .5
        if p > 1/2 then
            NeedToMove ← True
        end if
    end if
end for
```

there is a neighbouring location with a higher value of k_i. Dissatisfied agents that cannot find a neighbouring tile with a higher value of k_i compared to their current location, will remain stationary.

In the event that an agent will find itself isolated, with all neighbouring tiles vacant, the agent will always move. Isolated agents are always dissatisfied. To be in a norm-minority is preferable to being on their own. Agents of type 'Intelligent' and 'Risky' will consider the best vacant neighbouring tiles whose respective neighbourhood has the highest preferred norm plurality. 'Random' and 'Semi-intelligent' agents will chose a random vacant tile. If the neighbourhoods of neighbouring vacant tiles are also vacant (i.e. the agent would end up in isolation), all agents will move randomly.

The different reasoning patterns are designed to test the impact of flawed assumptions or false information. Because pluralistic ignorance relies on the sending of misleading signals, I test the impact that relying on these signals can have. The control treatment serves as a means of distinguishing patterns from noise: the occurrence of pluralistic ignorance through random movement serves as a baseline noise that any treatment is compared to. In effect, the different treatments serve as a

Algorithm 4.3 Move intelligently

```
if NeedToMove = True then                                         ▷ See alg. 4.2
    k⋆ ← highestNumberOfNormEquals
    L⋆ ← bestFutureLocation
    L_c ← SurroundingEmptyNeighbourSites    ▷ surrounding locations in range
nbr
    for all l ∈ L_c do
        s ← |{a_j ∈ N_t(l) : i_i = p_j}| + 1   ▷ number of nearby public norm equals
        d ← |{a_j ∈ N_t(l) : i_i ≠ p_j}|       ▷ number of nearby public norm differents
        k = s − d
        for all l ∈ L_c do
            k⋆ ← {k ∈ K_t(l) : k⋆ ≥ k ∈ K_t(l)}       ▷ find the highest possible k
            if k(l) = k⋆ then
                L⋆ ← L⋆ ∪ {l}                      ▷ update satisfactory locations
            end if
        end for
    end for
    l⋆ ← ChooseOneAtRandom(L∗)
    p_{i,t+1} = l⋆                                          ▷ update location
end if
```

proxy for reliance on information (signals from other agents). The *intelligent* agents that reason through both of their decisions rely on information the most: they trust the public signals of their fellow agents both when they reason if they should move, and if so, where they should go. The *semi-intelligent* and *risky* treatments each rely on public signals only during one of their reasoning steps: *risky* agents consider signals when deciding where to move, and *semi-intelligent* agents consider them when deciding if they move. The two treatments thus halve their reliance on public signals, yet at different stages. This serves as a way to determine what the impact of movement is: as movement is crucial in Schelling models, and agents are dependent on it to reach states of happiness. *Semi-intelligent* agents are able to ascertain whether they are indeed happy or not, *risky* agents are not. Any differences between the two treatments will thus be down to the probability of moving, and the information considered at the different stages of reasoning.

Agents cannot switch behaviours or norm adherence. In reality, people are able to change their attitudes and behaviours. The focus of this model is to explain misconceptions of attitudes- not their change. I explain the emergence of pluralistic ignorance on the basis of the desire to belong to a social group and the possibility of movement. There is no pressure exerted other than the mere presence of others. Agents can only escape or enter pluralistic ignorance through movement. Movement through social space is not equivalent to movement through geographical space, which can be costly. Social environments can be altered by reducing or increasing social interactions with different factions of friends, relatives, colleagues; or by changing the people and pages to follow on social networks (McPherson et al., 2001). The structure of space is another possible addition to the model: if the world wasn't a toroid (wrap-around, as this one is), agents in corners and on the side of the board would have a restricted neighbourhood. This would alter their reaction based on the neighbourhood ratios k and potentially result in border-agents that behave differently. A possible effect could be that border- or corner-agents are more erratic, as the reduced number of neighbours leads to fewer scenarios that are satisfactory: it is easier to satisfy the condition of 10/20 agents compared to 1/2. A single moving agent has a larger effect if the neighbourhood is so small. Because the model is held closely to the Schelling model and other ABMs that have modelled the spread of norms using Schelling, this feature is omitted. This has the benefit of comparing the model outcomes to similar models whilst making fewer changes at once. However, the concept of a restricted neighbourhood for corner- or border-agents could be a useful feature to simulate real-world situations in which people are socially cornered: for example, poverty could be a reason why social life is restricted to family or family and work- people can't afford to expand their social space.

4.4.3 Initial conditions

At the start of each simulation, initial conditions are randomised to test the robustness of the model.

1. **Population density**

   The population density $PopDen$ determines the number of agents relative to the grid size, $S \times S$. At the beginning of each simulation, the population density is drawn randomly from a uniform distribution $\sim U(0.25, 0.98)$. Typically, Schelling models are run under conditions of high density. Because agents in this model have limited movement and the 2-dimensional space represents a social environment, low densities such as 25% are also considered, representing a small social environment. The reduced movement capability lessens the impact that density usually has. Higher densities restrict freedom of movement, but agents in this case are restricted to the vacant tiles of their local neighbourhood.

2. **Neighbourhood range** The neighbourhood has a range *nbr* which specifies how far the agent's neighbourhood reaches. Figure 4.3 shows a visualisation of the different neighbourhood ranges. The smallest neighbourhood is a *von Neumann* neighbourhood with 12 tiles in a diamond shape. The medium and large neighbourhoods are Moore neighbourhoods with 24 and 48 tiles respectively. The different neighbourhoods were chosen so that each range doubles the number of tiles, making it easier to track any impact that the neighbourhood size has. The range *nbr* of $1 - 3$ refers to the radius of the neighbourhood.[9]

   Agents with $nbr = 3$ have potentially 48 tiles to move to (if vacant). A larger range thus increases the freedom of choosing a better area. More tiles also create more variations of s to d ratios.

   Neighbourhood size is usually constant for Schelling models and should be kept so to ease model comparison (see Section 2.2.2). However, because the grid in this model is meant to represent a far smaller environment, agents have

[9]Because the von Neumann neighbourhood is diamond-shaped, its range is not even.

Figure 4.3: Agent neighbourhood ranges

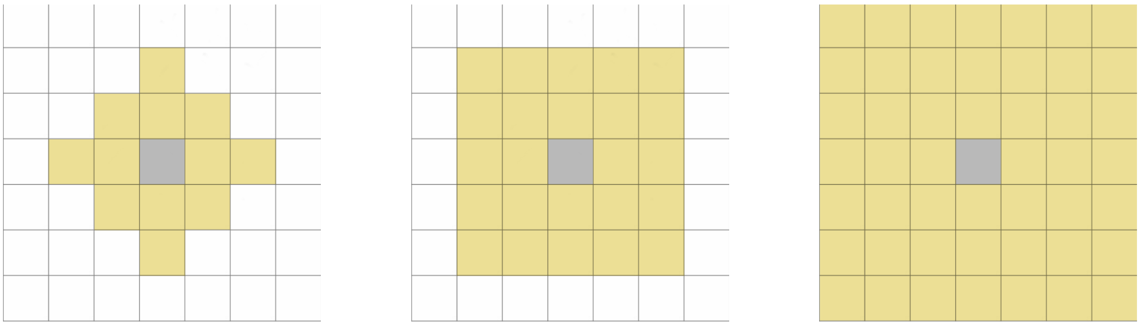

(a) Range $nbr = 1$, 12 tiles

(b) Range $nbr = 2$, 24 tiles

(c) Range $nbr = 3$, 48 tiles

limited movement and cannot 'jump' across the grid. Because their choices of movement are restricted to their local neighbourhood, the neighbourhood size acts as a proxy for freedom of movement. A neighbourhood of 12 compared to 48 tiles at the same population density is more likely to offer fewer choices. If 50% of neighbourhoods are occupied, there would still be 24 empty tiles to choose from, as opposed to just 6. This could influence the opportunity for agents to maximise their utility, and thus I include the neighbourhood range as a parameter.

3. **Number of norm-consistent agents**

   The number of agents that are norm-consistent is given by NC, drawn randomly from a uniform distribution $\sim U(0.2, 0.8)$. NC rates represent a range of minority, equality and majority situations for agents that follow either AA or BB. Each group (consistent and inconsistent agents) consists of a 50:50 split between the norm combination. If $NC = .25$, then .12,5 of all agents are AA, .12,5 BB, .37,5 AB and .37,5 BA. This split persists throughout every treatment of NC.

Table 4.3: Independent variables

Parameter	Range	Description
$MoveRule$	$\in Random, Intelligent, Semi-int., Risky$	Movement rule
$PopDen$	$\sim U(0.25, 0.98)$	Final population density
NC	$\sim U(0.2, 0.8)$	No. of norm-consistent agents
nbr	$(1-3)$	Neighbourhood range

4.4.4 Dependent variables

Each simulation runs for $t_{max} = 10,000$ ticks. Every 10 ticks, the dependent variables are sampled, allowing for time-series analysis of the outcome variables. Each treatment was repeated $3,000$ times, resulting in $4 \times 3,000 = 12,000$ cross-sectional samples and 1.2×10^8 time-series samples of the dependent variables. These are described below and summarised in Table 4.4. The end-state values of each dependent variable is denoted by removing the t subscript.

Table 4.4: Dependent variables

Variable	Description
M_t^Ψ	Segregation of Public Norms at time t (equation 4.5)
M_t^Ω	Segregation of Private Norms at time t (equation 4.6)
$\bar{O}_t$	Number of occurrences of PI at time t (alg. 4.4)
$\bar{Z}_t$	Average size of PI groups at time t (alg. 4.4)
IC_t	No. of norm-inconsistent agents in PI groups at time t
$\bar{v}_t$	average agent movement at time t
M^Ψ	Segregation of Public Norms at the end of the simulation
M^Ω	Segregation of Private Norms at the end of the simulation
IC	No. of norm-inconsistent agents in PI groups at the end of the simulation
$\bar{O}$	Number of occurrences of PI at the end of the simulation
$\bar{Z}$	Average size of PI groups at the end of the simulation

I record the Moran's index of spatial autocorrelation for both public norms and private norms in order to quantify the segregation that exists between the norm groups. Because agents consider their own private norm when comparing it to others' public norms, the two segregation indexes are correlated. The segregation of public norms is given by:

$$M_t^{\Psi} = \frac{|A|}{\sum_{(i,j)\in A^2} w_{i,j}} \frac{\sum_{(i,j)\in A^2} w_{i,j}(\Psi_i - \bar{\Psi}_t)(\Psi_j - \bar{\Psi}_t)}{\sum_{i\in A}(\Psi_i - \bar{\Psi}_t)^2} \tag{4.5}$$

where the mean public norm is $\bar{\Psi} = \sum_{i\in A_t} \Psi_i/|A_t|$, and $w_{i,j} = 1$ iff agents a_i and a_j are immediately adjacent on the grid (including diagonals), otherwise $w_{i,j} = 0$. The equivalent segregation index for private norms is:

$$M_t^{\Omega} = \frac{|A|}{\sum_{(i,j)\in A^2} w_{i,j}} \frac{\sum_{(i,j)\in A^2} w_{i,j}(\Omega_i - \bar{\Omega}_t)(\Omega_j - \bar{\Omega}_t)}{\sum_{i\in A}(\Omega_i - \bar{\Omega}_t)^2} \tag{4.6}$$

Because the segregation indexes measure public and private norms independently, I record the number of Pluralistic Ignorance groups (PI groups), $\bar{O}_t$ and the average size of all groups $\bar{Z}_t$ using an implementation of a *FloodFill* algorithm 4.4. Let L_A be the positions of all agents A_t and all tiles W of the grid. The world is searched for a given public norm $\Psi = A$ or $\Psi = B$. A tile with an agent of the given public norm **and** an inconsistent private norm is added to the list of PI groups. If adjacent agents show the same norm inconsistency, they are added to the same group.

Algorithm 4.4 FloodFill

```
function FloodFill(z,Ψ)
    n ← Ψ = A                                  ▷ the norm to be searched
    F ← {}                                     ▷ initialize flagged tiles
    for w ∈ W : w ∉ F do
        if w ∈ L_A then                        ▷ If tile has agent
            if Ψ(a_i) = n ∧ Ω(a_i) ≠ n then    ▷ If the agent is norm inconsistent
                F ∪ w                          ▷ add tile to flagged tiles
                FloodFill(w, Ψ)
            end if
        end if
    end for
end function
```

The *FloodFill* thus captures either AB agents or BA agents. AA and BB agents are not considered. Norm consistent agents are considered when inconsistent agents move. Groups with fewer than three agents are not considered PI groups. The minimum of three was chosen because PI definitions require a plurality of norms, which is not possible with two agents (which would tie) or one agent, which does not constitute a group.

To capture prevalence of pluralistic ignorance, I record the number of norm-inconsistent agents that are in PI groups, IC_t. Because the number of inconsistent agents is predetermined by NC, this parameter serves as a measure of how many inconsistent agents manage to escape pluralistic ignorance. As groups of PI are defined by a minimum of three adjacent agents with the same norm inconsistency, some PI groups can occur by chance (through random movement). The measure of IC_t can be compared between the treatments to describe propensity of PI emergence.

$$IC_t = \frac{\bar{Z}_t \times \bar{O}_t}{|a_i \in A : \Psi_i \neq \Omega_i|} \tag{4.7}$$

Lastly, I capture average movement of agents, capturing the average number of agents that have moved at time t:

$$\bar{v} = \frac{|a_i \in A : p_{i,t} \neq p_{i,t-1}|}{|A|} \tag{4.8}$$

Table 4.5: State variables

Variable	Description
A	The population of agents
Ψ_i	Public Norm of agent a_i
Ω_i	Private Norm of agent a_i
$p_{i,t}$	Position of agent a_i at time t
P_A	The positions of all agents
$N_t(a_i)$	The set of agents that are neighbours of agent a_i
$N(p)$	The set of locations in the neighbourhood of location p

4.5 Results

4.5.1 Typical simulation runs

Figure 4.4 visualises the clustering arrangements of private norms Ω across the different movement treatments and over time. Ticks $t = 1, 10, 100, 500$ are presented (left to right, each row).

The random movement(a) shows no clustering. Intelligent movement (b) shows clusters typical of a Schelling model: because agents seek others with the same norm convictions, they end up forming visible clusters that grow larger over time. Because NC is 50% for these model runs, only half of the visible private norms will be consistent with what the public norm signals.

Semi-intelligent movement (which includes 'whether-to-move' reasoning, but not 'where-to-move') shows clustering as well, but less dense. Agents broadly align into clusters, but empty space is scattered across and between the different norm groups.

Risky movement shows the least jagged edges: agents form one large cluster that is not very segregated (many areas contain both norms) but has very smooth edges.

4.5.2 Quantitative analysis

The emergence and occurrence of pluralistic ignorance is tested by comparing different movement rules and varying the model parameters described in Section 4.4. Figure 4.5 shows a scatterplot of IC, the fraction of possible PI agents that is part of PI groups against the norm consistency of agents, NC. The figure also distinguishes between the different neighbourhood ranges, nbr. As a general rule, the more norm-consistent agents exist, the lower is the overall propensity of agents to get caught in pluralistic ignorance situations. This is intuitive: more consistency between private and public norms means fewer misleading signals that other agents pick up and thus, the probability that agents experience pluralistic ignorance goes

Figure 4.4: Comparison of cluster arrangements of Ω between Movement conditions. Parameters: $PopDen = 62$, $nbr = 3$, $NC = 50$. Time $t = 1, 10, 100, 500$.

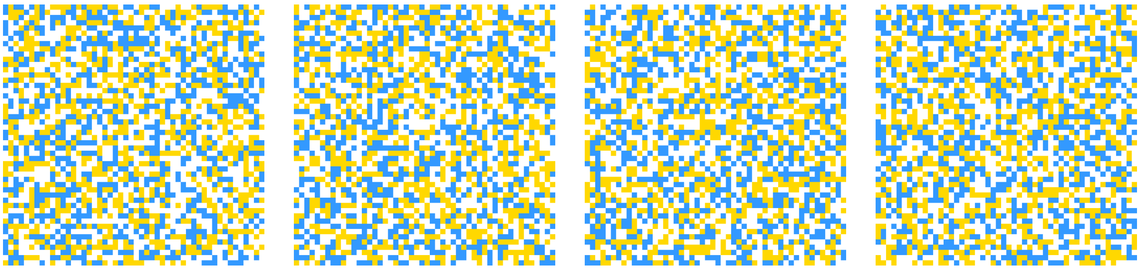

(a) Random movement

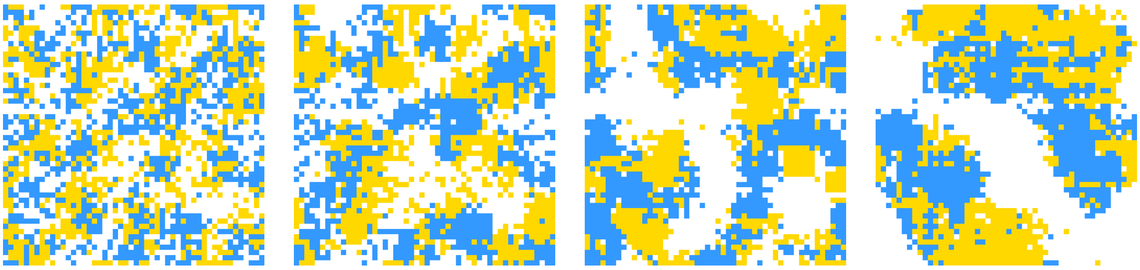

(b) Intelligent movement

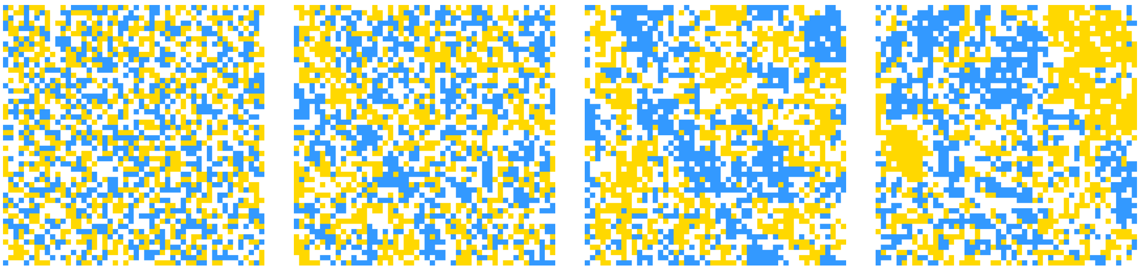

(c) Semi-intelligent movement

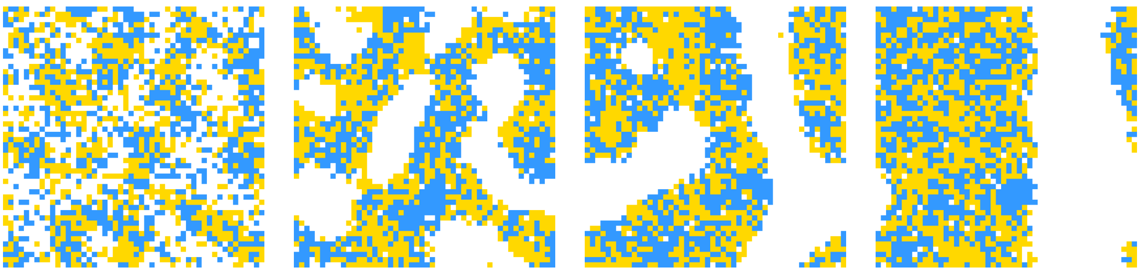

(d) Risky movement

down (misrepresentation of attitudes is a crucial element of PI, as discussed earlier). Higher neighbourhood ranges affect *intelligent* and *risky* agents the most, those are the agents that reason when deciding where to move (Table 4.2). *Intelligent* agents have higher PI rates when *nbr* increases, indicating that their reasoning fosters the advent of PI when the range of options is greater. This is due to the intelligent agents' desire to prevent isolation and seek friends. The farther they can see, the more options they have to find friends and, ultimately, fall prey to faulty information. The worst performers, judged by how much PI they suffer, are *risky* agents with $nbr = 3$. They experience much higher levels of PI when NC is between 65% and 70% (Figure 4.5: l). The key to these differences is population density.

Density of the population, *PopDen*, affects the number of agents that form part of PI groups. Higher levels of *PopDen* are denoted by darker colours of the grey-scale. Except when agents are *risky* (j,k,l), lower levels of *PopDen* generally correlate with lower values of IC. *Intelligent* agents with larger *nbr* ranges (e,f) don't benefit from this effect as much: similar to their *risky* counterparts, the high levels of IC continue as NC increases and only drop off when NC passes 60%.

The risky agents' PI results show far less spread because density no longer impacts the results as it does with random, semi-intelligent and intelligent agents. Why is that? The effect of density is to provide more room for movement when density is low, and that also means that there are more ways to be satisfied: in high density areas, agents have fewer choices to move to, even if they want to. This perpetuates PI situations. Risky agents however ignore the local population density when they decide whether to move, because their reasoning is random. That means that when a situation is satisfactory, the agent still has a 50% chance of moving away. A low-population situation in which all agents stop moving because all of them have found a satisfactory place is not possible. The movement indicator (red line) highlights the high numbers of movement. Contrary to random agents, who also move by chance, risky agents will seek satisfactory locations once they have to move. Finding new places based on faulty information creates PI situations, and lower densities alleviate this effect. Risky agents cannot benefit from this mechanic.

When NC is between 60% and 70% and $nbr = 3$, *risky* agents experience a

Figure 4.5: IC by NC, controlling for $moveRule$, $popDen$ and nbr. The x-axis shows the proportion of norm-consistent agents (NC). The y-axis shows the amount of norm-inconsistent agents in PI, IC. The z-axis is depicted as a gradient, showing the different population densities between 25% and 98%. Darker shades of grey denote higher densities. Each row shows a different $moveRule$, each column shows a different neighbourhood range nbr. The average agent movement $\bar{v}$ is depicted by the red lines and share the scale of the y-axis of 0.0 to 1.0.

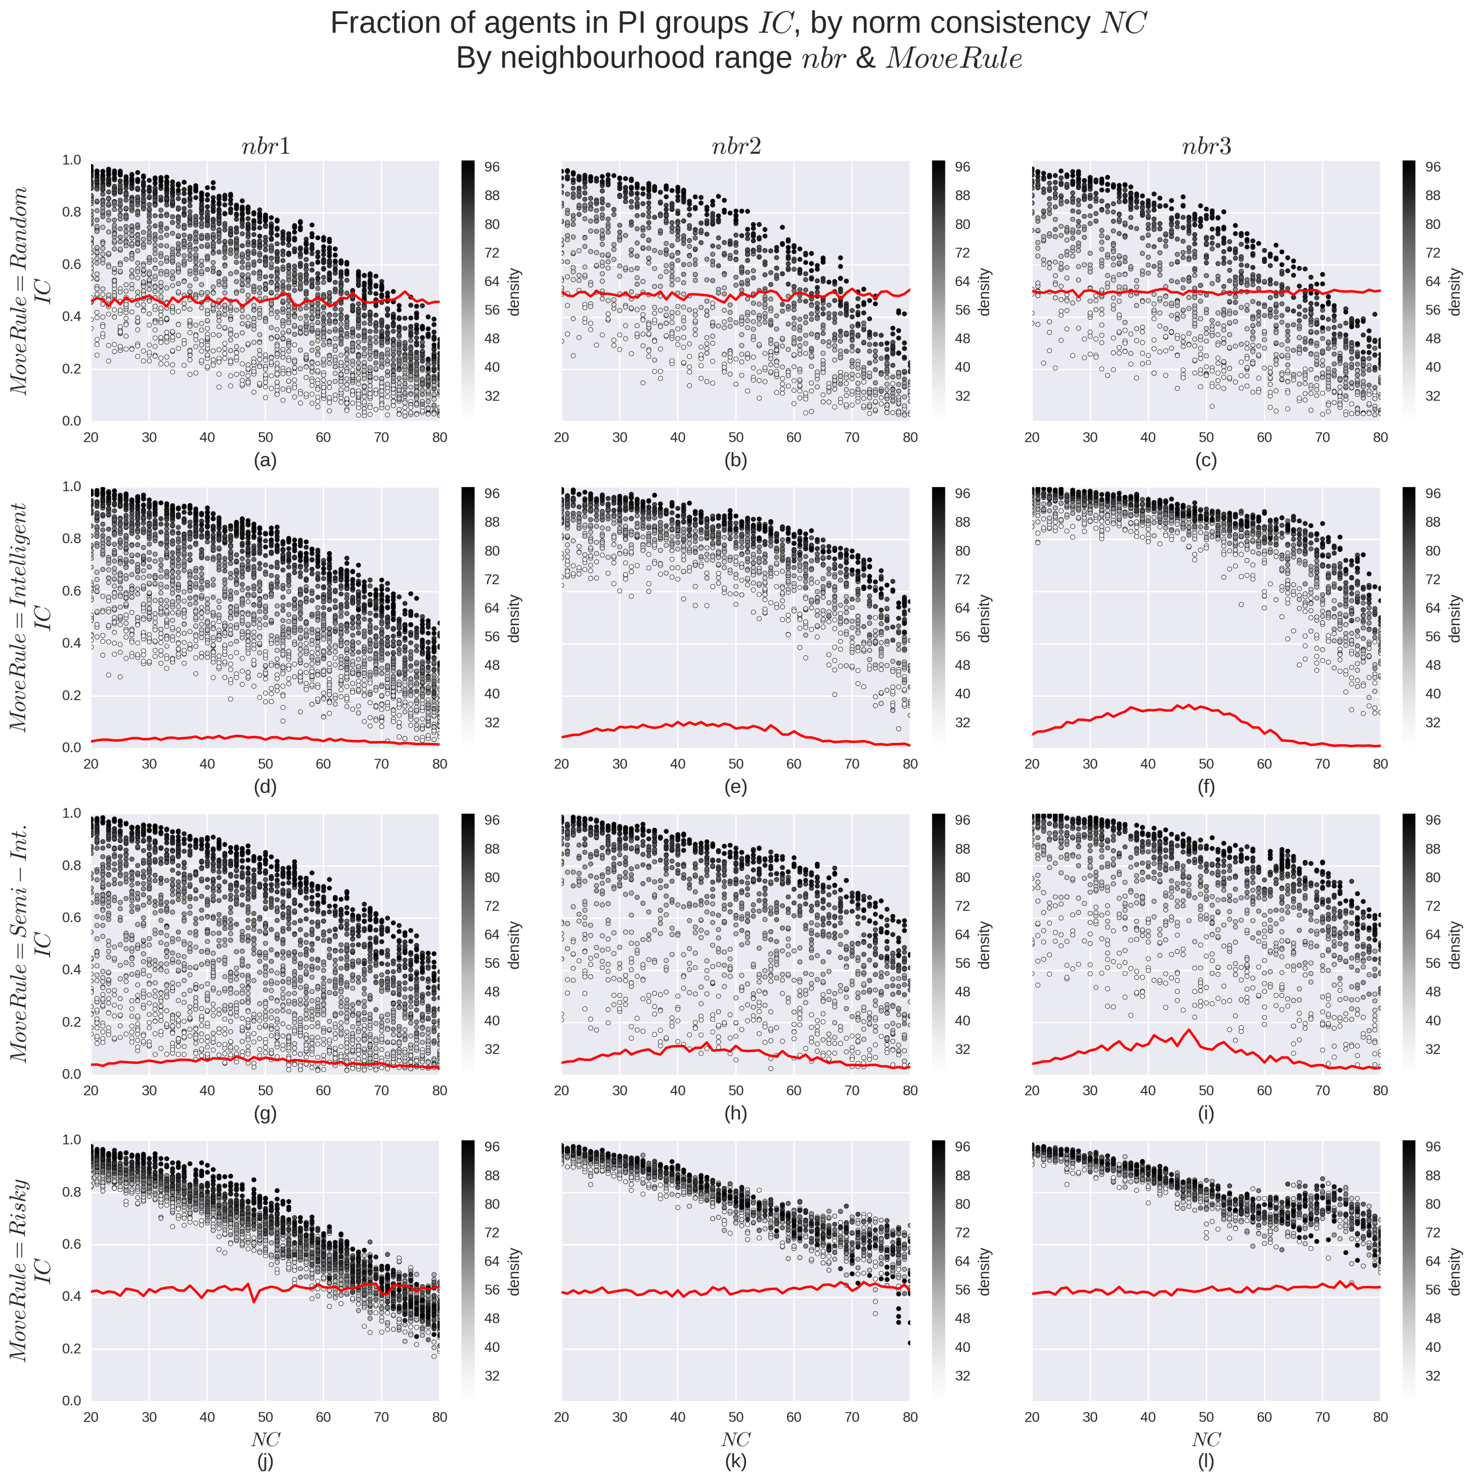

further uptick in PI (l). The same effect is hinted at when $nbr = 2$ (k), but the variance is too large to make out the same effect clearly. Contrast this to intelligent agents at $nbr = 3$ (f): once NC levels of 60% are reached, PI drops sharply. This is the point at which norm-consistent agents form a larger majority and the chances of being surrounded with genuine agents are higher. Intelligent agents stop moving, as do semi-intelligent agents: their movement indicator drops to near-zero at NC=70%. More consistency leads to less of a domino-effect: with norm-inconsistent agents, a new neighbour who thinks they have found the right place tips off existing neighbours who have been untruthful and move away as a result. This is less likely with higher levels of NC. Thus, movement grinds to a near-halt and many agents reside satisfied in places that, depending on the population density, have high PI (high *popDen*) or low PI (low *popDen*). Risky agents, through their push to move, effectively raise the required number of norm-consistent agents for PI to go down. Because they keep moving and prevent non-PI happiness, the majority of NC has to be so overwhelming that PI can be reduced. Their reliance on signals of others seals their fate, as it were.

As indicated above, the different reasoning patterns of agents affects their average movement (red lines). As expected, *random* and *risky* agents have movement levels close to .5 and .4 respectively. This 10% difference is due to their different neighbourhood considerations: *random* agents check their neighbours for empty, *risky* agents check their neighbours' neighbours as well as the suitability of new places, and can change their mind. When $nbr = 1$, *intelligent* agents achieve the same results (d) as their *random* counterparts (a) with far less movement: across all values of *nbr*, both *intelligent* and $semi-intelligent$ agents move very little (d-i).

Greater neighbourhood ranges results in these agents moving more, with an average of 20% of agents moving at times (f,i). This is due to the greater range of locations to move to. With only 12 neighbours, it's less likely to find an improved empty location compared to 48 neighbours. When the range of neighbours increases, the numbers required to meet or miss a preference target changes. I.e. out of four neighbours, only two friends are required for a majority, but out of ten agents, at least five need to be friends. The ratio doesn't change, but the number of possibilities to achieve that ratio increases. The neighbourhood range affects one element of agents' reasoning in particular: their drive to be in community with others, and not isolated. The farther an agent can see, the more options they have to overcome isolation quickly. We can see this effect looking back at 4.4: risky and intelligent agents, the two that reason where to go, show tighter clusters and clear edges, with large areas of empty space. Semi-intelligent agents who consider whether to move, but not where to, show a different pattern: they form clusters, but these clusters have ample empty room inside, leaving no discernible borders between different clusters. The density inside of the clusters increases under risky and intelligent movement, but not under semi-intelligent (and random) movement.

The overall 'winners' of keeping PI low are semi-intelligent agents, particularly those when $nbr = 1$ (Figure 4.5 g). They outperform random movement especially when NC is very low. In low-density situations, IC is often lower than 0.2, i.e. less than 20% of all agents that could be in a PI situation, are in PI situation. When misinformation is rife, ignorance is bliss: because semi-intelligent agents leave it to chance where to move to, they don't suffer the effects of relying on wrong signals (intelligent agents do). And because they stop moving once they are happy, they allow neighbourhoods of low PI to settle and don't generate new chances of PI (risky agents do).

Figure 4.6 shows the levels of IC at the end of simulations as boxplot diagrams. Norm consistency and population density are grouped into four and three different ranges respectively to aid comparison. When density is low (a-d), truthfulness of agents reduces pluralistic ignorance. When more agents are truthful, inconsistent agents can find satisfactory groups made of consistent agents and avoid PI. When keeping truthfulness constant, the movement rule accounts for major differences in PI propensity. When truthfulness is low (a,b), $semi-intelligent$ agents outperform $random$ agents in avoiding pluralistic ignorance. When truthfulness is above 50% (c,d), $random$ agents marginally outperform $semi-intelligent$ agents. $Intelligent$ agents suffer up to 30% more agents in pluralistic ignorance compared to $random$ movers, but $risky$ agents have the largest PI occurrences throughout the range of NC (a-d), for the reasons discussed above.

For medium $PopDen$ levels (e-h), the movement treatment differences follow the same pattern, but pluralistic ignorance is higher overall. It is no longer the case that $semi-intelligent$ agents outperform $random$ movers, but they remain the best-performing of the reasoning agents. The differences between $intelligent$ and $semi-intelligent$ behaviour are reduced. Consistent with intuition, overall, higher levels of truthfulness still drive down levels of IC as more truthful signalling will reduce the number of chances of PI occurring in the first place.

When density is high (i-l), the number of agents in pluralistic ignorance rises as well. A crucial difference is that $risky$ agents are no longer the worst performers. As truthfulness goes up (k,l), they out-perform both $intelligent$ and $semi-intelligent$ agents, but $random$ movers enjoy lowest values of IC. The combination of high density and low truthfulness (i,j) features the highest levels of PI: nearly all of the agents that are norm-inconsistent are in a pluralistic ignorance situation. The different movement rules can no longer mitigate the occurrence of PI, as they do under conditions of medium truthfulness and low population density. This shows that reasoning agents can indeed lower the propensity of PI, but that their abilities are

Figure 4.6: Propensity of norm-inconsistent agents to be part of PI groups

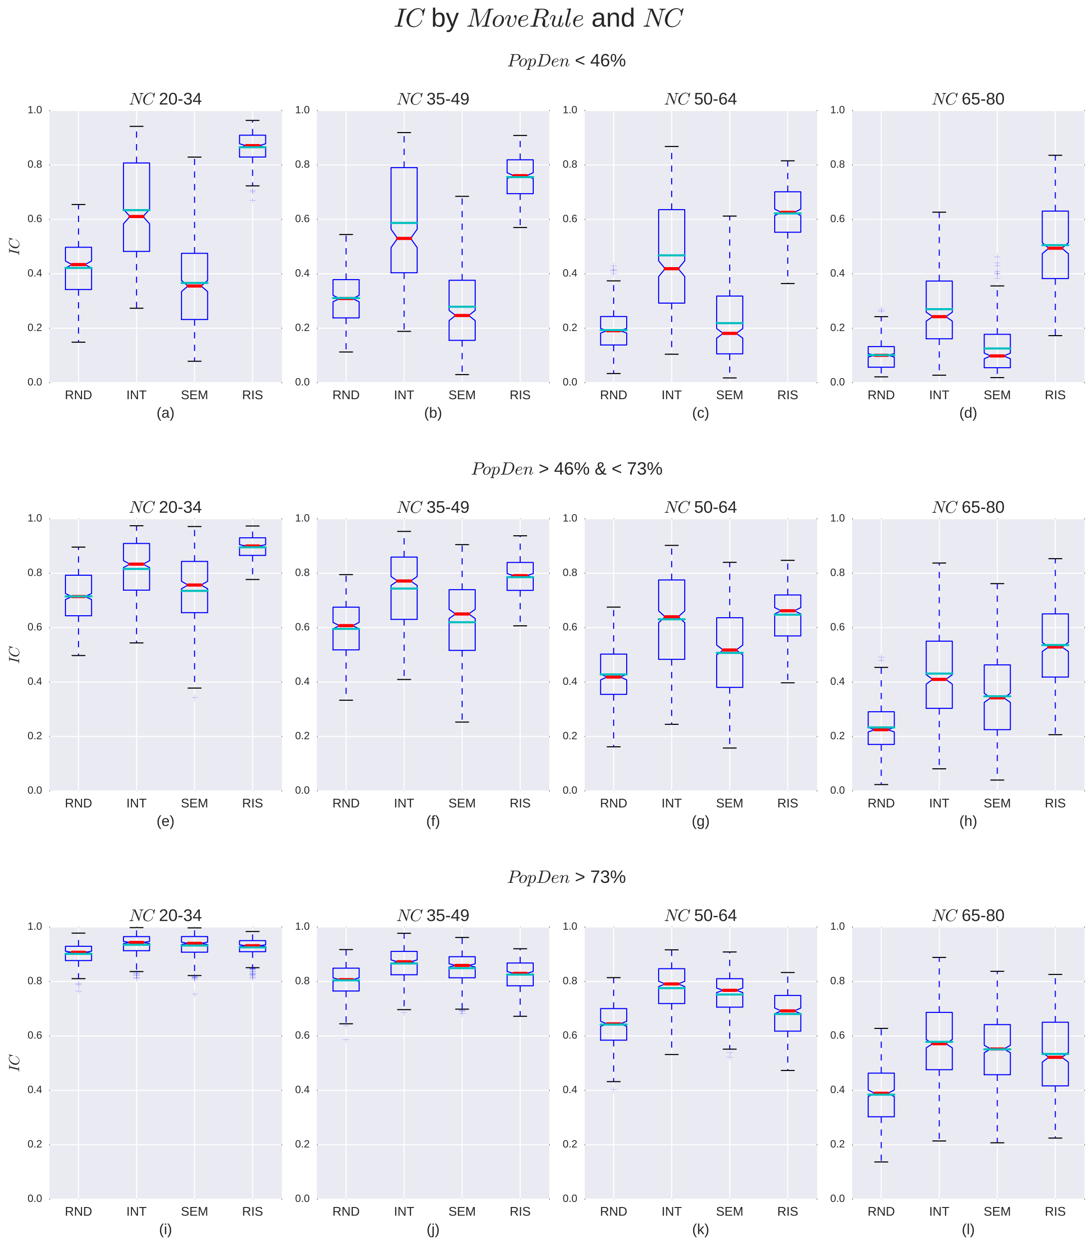

restricted to low and medium-range values of *popDen* and *NC*. The higher the population density, the less pronounced are the movement rule effects and differences. This affects $semi - intelligent$ agents the most. This is because the higher density increases the chances of inconsistent agents when agents survey their neighbourhood. The more deception there is, the more likely it is that reasoning agents draw the wrong conclusions. For the $semi - intelligent$ agents, the random movement away is no longer an advantage when the world is densely populated. Further evidence that their strategy stops working is visible in 4.6 i, j: the levels of IC of $semi-intelligent$ and $intelligent$ agents are nearly identical, in both levels and pattern. Only as consistent agent numbers increase further, can $semi-intelligent$ agents gain some ground compared to their $intelligent$ versions, but the differences are very small.

The IC variable is computed using average PI group sizes $\bar{Z}$ and numbers of groups $\bar{O}$ (see equation 4.7). Figures 4.7 and 4.8 show boxplot diagrams of mean group size $\bar{Z}$ and number of groups $\bar{O}$ instead of the computed value of IC. Again, the diagrams are broken down by $moveRule$ and NC, and split into low, medium and high population densities $popDen$. This can show us which patterns of IC in general are driven by each of its components.

The relationship between higher NC and lower IC is driven by both the number of PI groups (Fig. 4.8) and their average size (Figure 4.7). The diagrams on each row show the same trend: as NC increases, $\bar{O}$ decreases. Similarly, in Figure 4.7, as NC increases, $\bar{Z}$ decreases, but the differences between $NC20-34$ and $NC35-49$ are more pronounced than any other. At low densities, there are very little changes. The average size of PI groups captures the differences between the movement rules more clearly. Mean group size of PI groups goes down as NC increases because the likelihood of inconsistent agents who find a PI situation decreases overall.

When density is high (Figure 4.7 i-l), the previously clear patterns of low group size and low variance for $random$ and $semi-intelligent$, and high group size and high variance for $intelligent$ and $risky$ agents, disappear. Measuring $\bar{Z}$ only, $risky$ agents outperform their reasoning counterparts at medium-high density (Figure 4.7 k). However the number of groups for the same scenario is not lower for $risky$ agents (Figure 4.8, k). This shows that the 'uptick' we saw in Figure 4.5 k, l, is driven by group size $\bar{Z}$. The diagram shows that the 'uptick' is a result of the variance in sizes: a large number of groups are larger than average, causing the spike in IC- the trend however, continues downwards.

When $NC < 34$, *risky* agents have the fewest groups, but large group sizes. This is due to the clustering pattern that *risky* agents show (see Figure 4.4 and the discussion above): their constant movement into densely populated clusters at the expense of free space that can break up PI clusters. When the density is higher, free space becomes limited regardless of the movement rule and stop affecting group sizes as much.

The mean group size is vastly increased when NC is very low (Fig. 4.7: a, e, i. Note: the y-axis is different for those three plots, ranging from 0-37 instead of 0-20). When density is high, the lowest truthfulness values (i) feature average groups of 10 agents more than the groups when NC is higher (j). The difference for low and medium densities (a,e) is less pronounced but still significant, with at least 5 agents more per group.$\bar{Z}$ is so high when NC is low because of the great majority of inconsistent agents. Especially when *popDen* is high and the map becomes crowded, even *random* agents cause large PI groups. This shows that the movement rules no longer account for $\bar{Z}$, but the higher quantity of agents and low number of consistent agents among them do. Mean group sizes grow so large when not enough empty space exists to break up existing groups. The floodfill algorithm considers adjacent tiles a group, regardless of how much of a cluster it might look like. A long line of agents constitutes a group as much as a cluster does. When NC is so low, the sheer quantity of inconsistent agents makes it more likely that, even by chance, a stray inconsistent agent provides a link between two existing PI clusters. The bad performance of *random* agents at higher densities (Figure 4.7 i) supports this: under these conditions, even completely random movement will result in very large PI groups. For densities less than 73%, *intelligent* and *risky* agents show similarities in mean group size. Their reasoning when picking a new place leads to larger groups- again, this can be captured visually by observing the lack of empty

Figure 4.7: Mean group size of PI groups

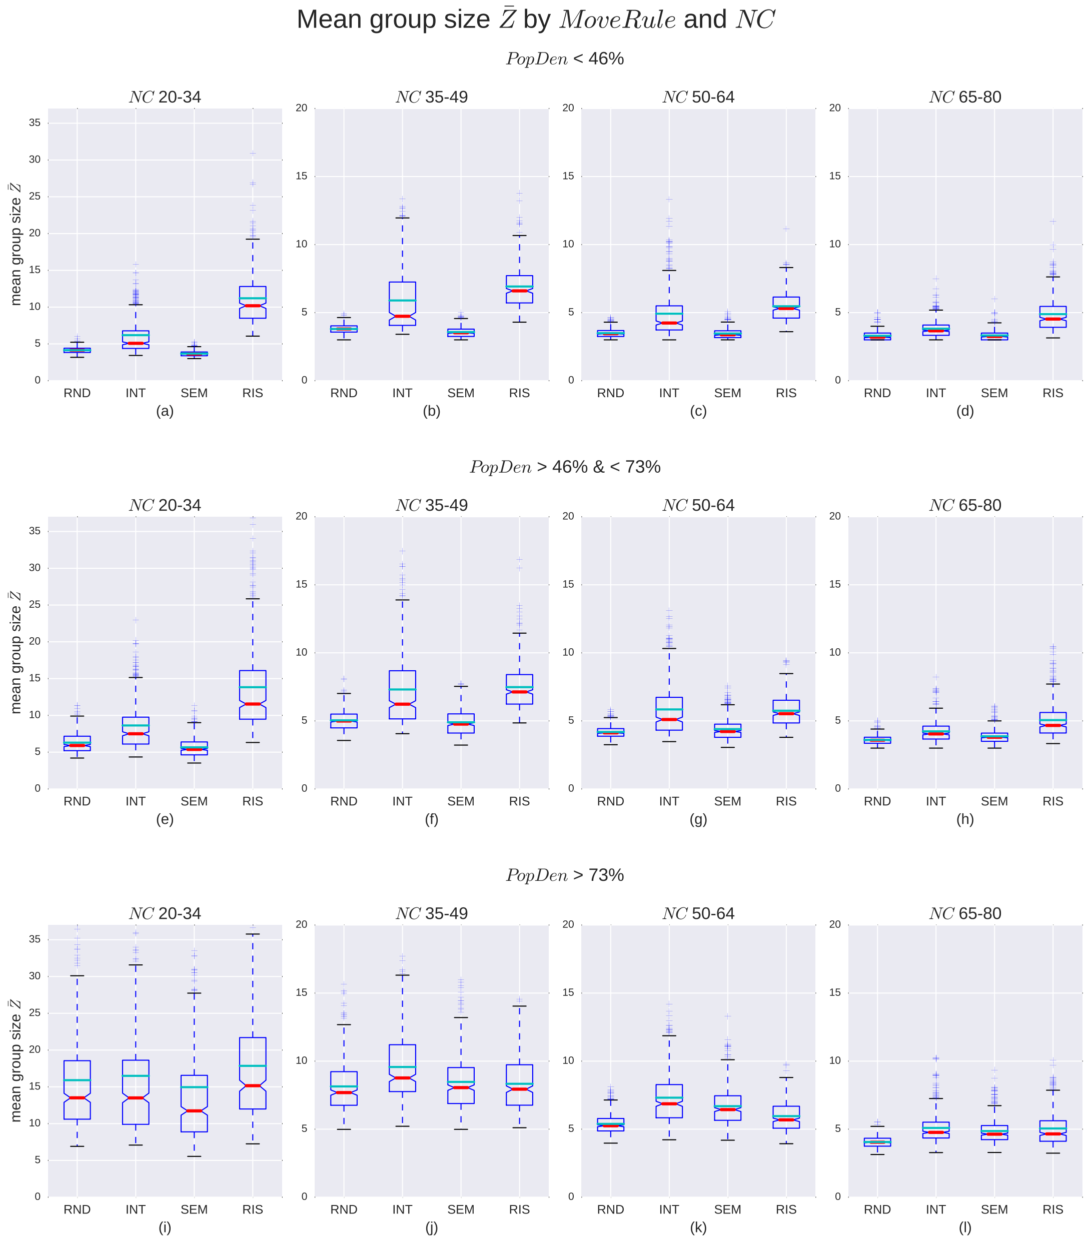

space during these runs. The mean group size boxplots confirm the importance of population density and their influence on the agents' movement behaviours, in line with the observed patterns of behaviour in Schelling models.

As described in Section 4.4.4, I also record segregation levels. Because agents essentially exist in two domains (public and private), I record each domain independently. Figure 4.9 shows average private norm segregation levels broken down

Figure 4.8: Number of occurrences of pluralistic ignorance groups

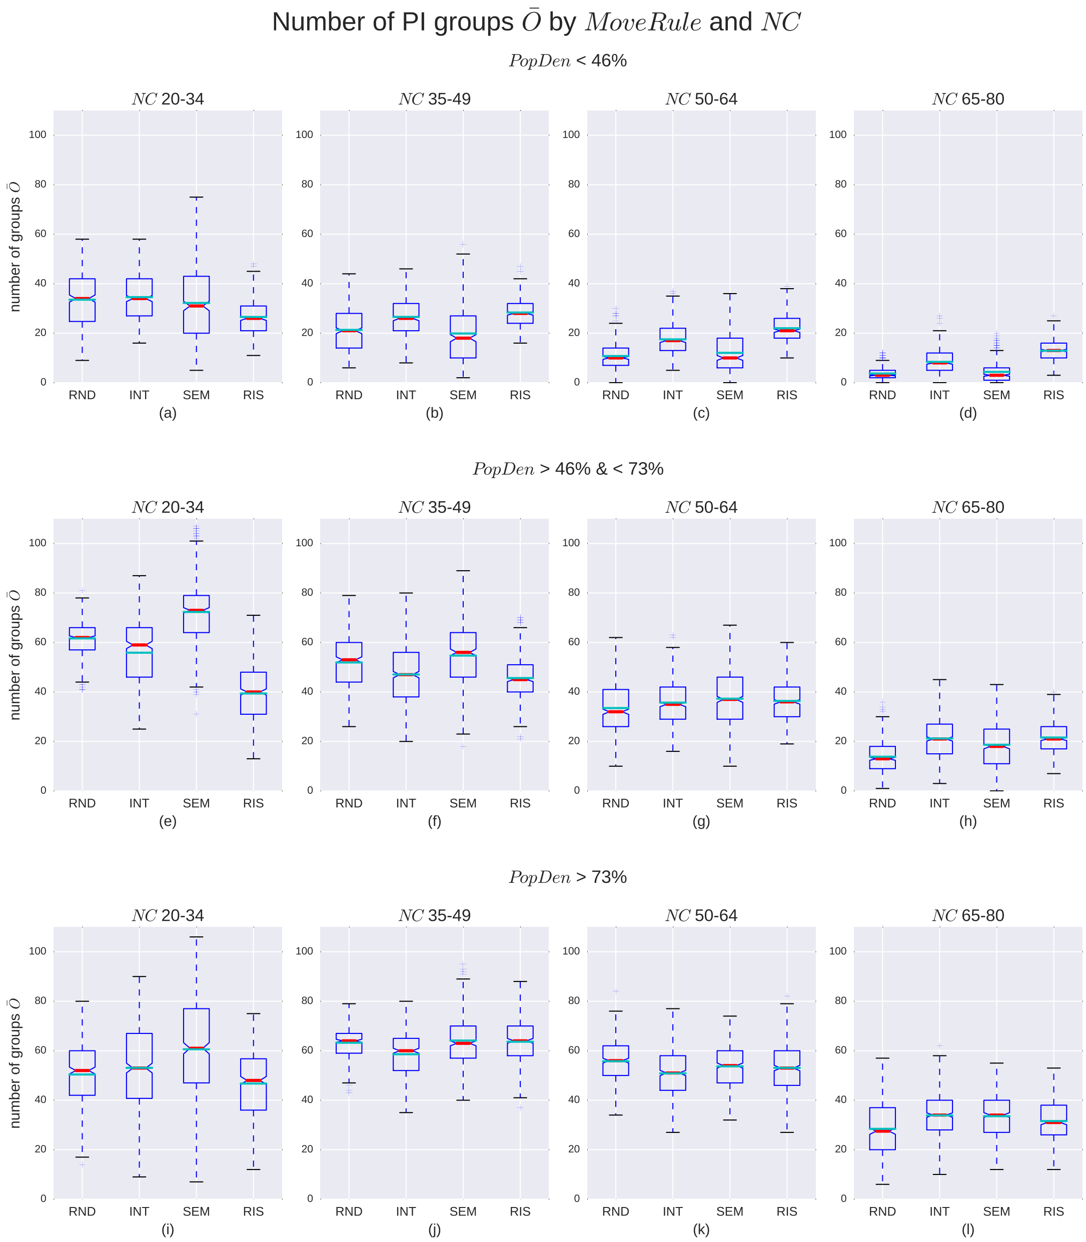

by $MoveRule$, nbr and $PopDen$ as previously shown in Figure 4.5. The columns represent the neighbourhood range nbr, the rows show the different movement rules. The x-axis is NC, and the z-axis shows density as a gradient. Private norm segregation is shown in shades of orange, public norm segregation in shades of blue. It should be noted that the *range* of segregation level differs in each figure: Figure 4.9 measures private segregation levels up to 0.8, Figure 4.10 scales not higher than 0.4.

Figure 4.9: Segregation of Private Norms M^{Ω}

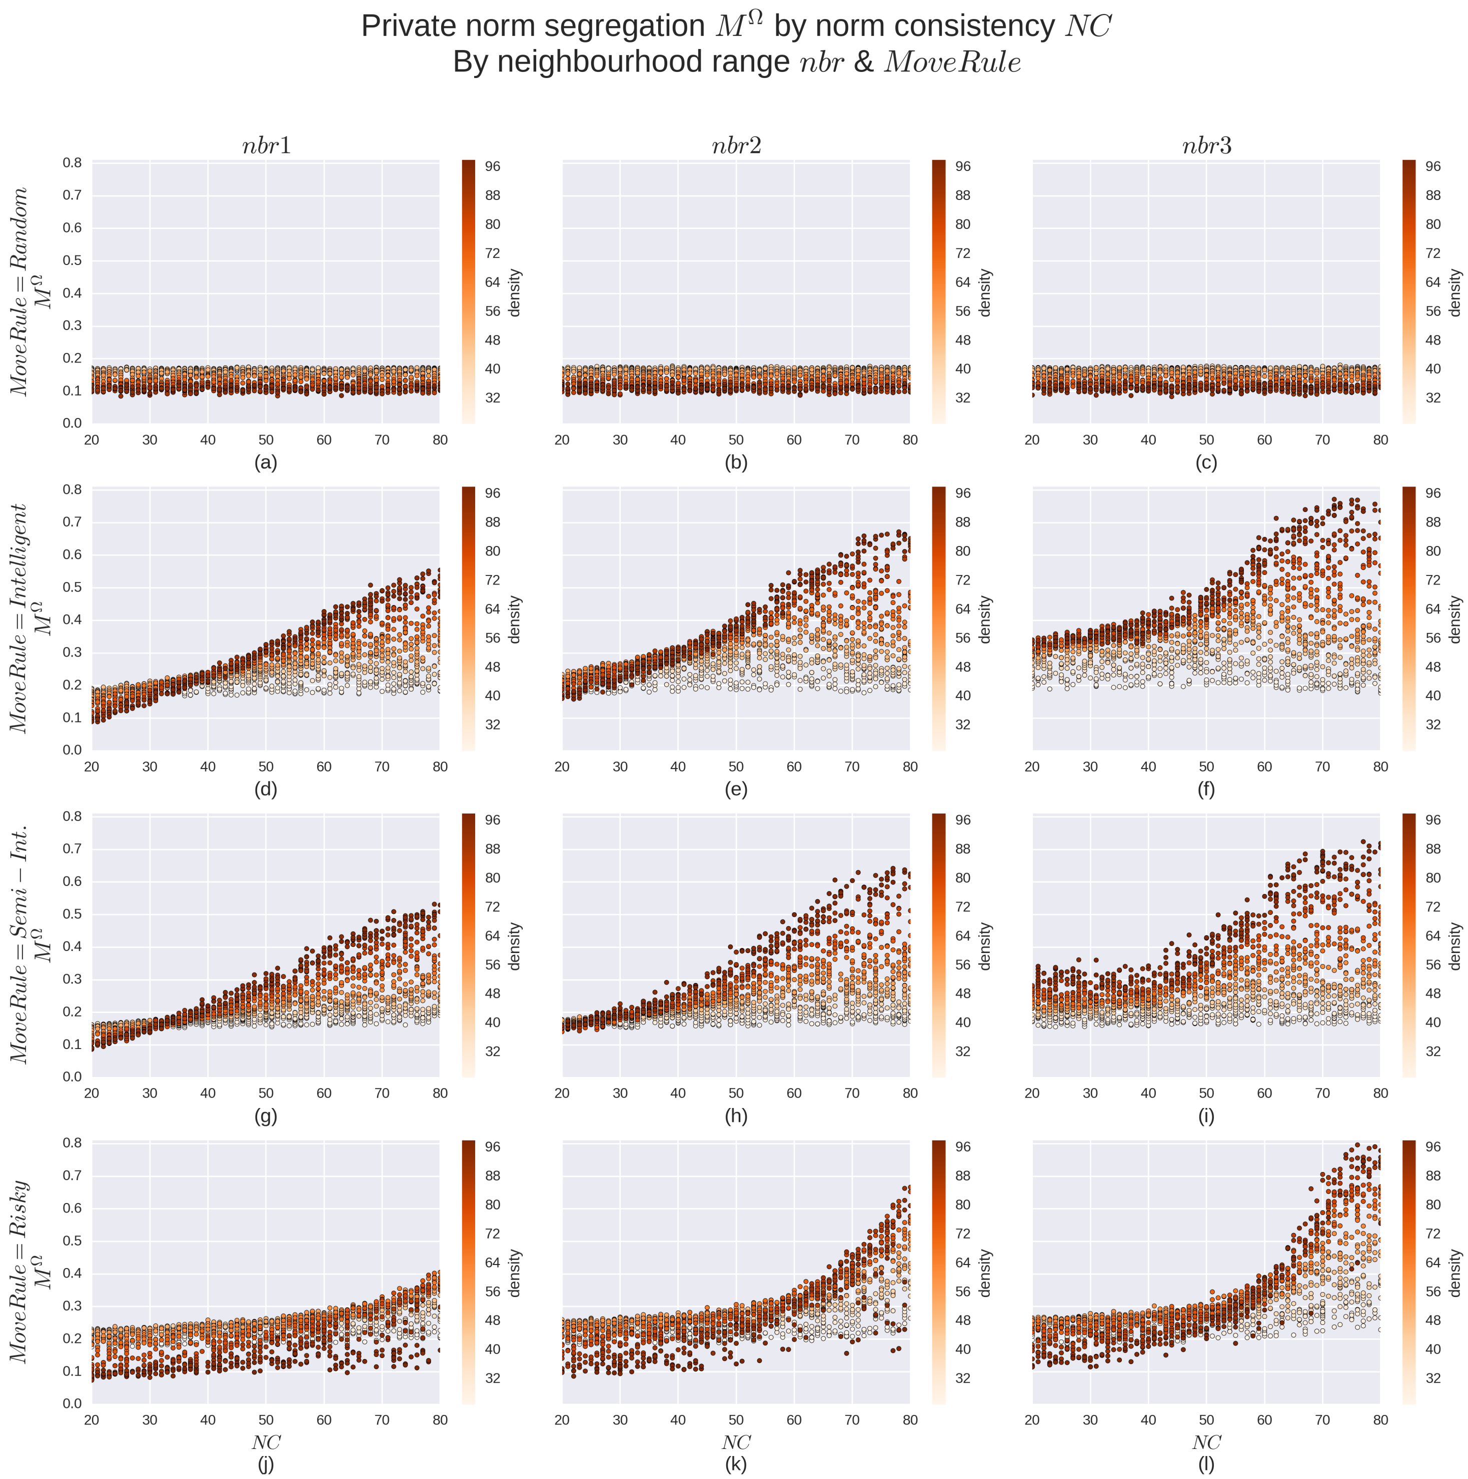

The control of random movement shows consistently low levels of segregation across all parameter ranges. The value is just below .2, as typical for random movement induced segregation in Schelling models. The three movement treatments affect segregation slightly differently. When agents move intelligently, segregation increases as NC increases, and density drives segregation higher still. This is true for all nbr treatments, but when $nbr = 3$, segregation is higher on average when NC is low. At peak levels of NC and $PopDen$, segregation levels are just shy of .8, whereas $nbr = 2$ peaks just under .7 and $nbr = 1$ at .54. The increased vision and equally, movement options of agents increases their ability to segregate according to their preferences (see the discussion on the effect of increased nbr above). Higher levels of nbr generally increase private norm segregation at all levels of density. Semi-intelligent agents show a very similar pattern of segregation compared to intelligent agents, showing that the whether-to-move reasoning is the driving factor of segregation. Agents stay in their neighbourhoods when these are satisfactory-which often means a level of segregation that agents are oblivious to. Risky agents on the other hand have lower segregation levels for a larger range of NC. When NC is less than 50, the segregation levels rise only marginally. Once NC reaches 60, private norm segregation shoots up, surpassing those levels found with *intelligent* and $semi-intelligent$ agents by the time $NC = 75$ is reached. The effect of $popDen$ is therefore delayed: only at $NC > 50$ does segregation begin to rise: the higher the nbr, the more rapid is the rise. At peak NC and $PopDen$ levels, only 40% agents with $nbr = 1$ are segregated, compared to nearly 80% of agents under $nbr = 3$.

When truthfulness increases, so does segregation. Because agents signal their private attitudes publicly, the decisions made where to move to and if to stay are largely based on true information. Wrong signals deceive agents to reside next to "false friends", thereby lowering segregation because agents are now mixed. Density increases segregation as truthfulness increases, because it increases the chances of responding to true signals. The movement rule effects again show a critical difference between randomly deciding whether to move, and randomly deciding where to move. *Risky* agents achieve low levels of segregation for a long period of NC because they don't stay where they are when satisfied. While this increases the occurrences of PI (see Figure 4.5 j, k, l), it does keep segregation down. Segregation levels are a result of the agents' behaviours, and are negatively correlated with pluralistic ignorance occurrences, IC.

Figure 4.10 shows the same breakdown of values, this time for *public* norm segregation M^{Ψ}. At first glance, patterns are similar to those of M^{Ω}: the control is consistently low, and the movement treatments show increases in segregation as NC and *PopDen* go up. However, public norm segregation is consistently lower than private norm segregation. The peak values, again when $nbr = 3$, $NC > 50$ and high *PopDen* levels, reach no more than .4, compared to .8 of private norm segregation. This is intuitive. Private norm segregation M^{Ω} should be higher than public norm segregation M^{Ψ}, because agents consider their own private norm when determining their happiness. When NC is low however, this effect is diminished: because such a large amount of agents send misleading signals, the population ends up mixed, as any attempt to bring own private norms together with other public norms will end up in a mismatch. Indeed, at very low levels of NC and $nbr = 1$, both M^{Ω} and M^{Ψ} are very low (both Figures, a, d, g, j).

Under intelligent movement and $nbr = 3$, high density and low NC values cause less segregation than lower density levels at the same parameter range, but quickly overtakes when NC approaches 65. Public norm segregation is also lower when density is high and agents are risky, as shown in the bottom row. Whereas M^{Ω} increases for risky agents as NC increases, this is not the case for public norm segregation M^{Ψ}.

When truthfulness is low and the majority of agents deceive with their publicly displayed attitudes, the public 'surface' will be a mixture of both norms, but the pri-

Figure 4.10: Segregation of Public Norms M^{Ψ}

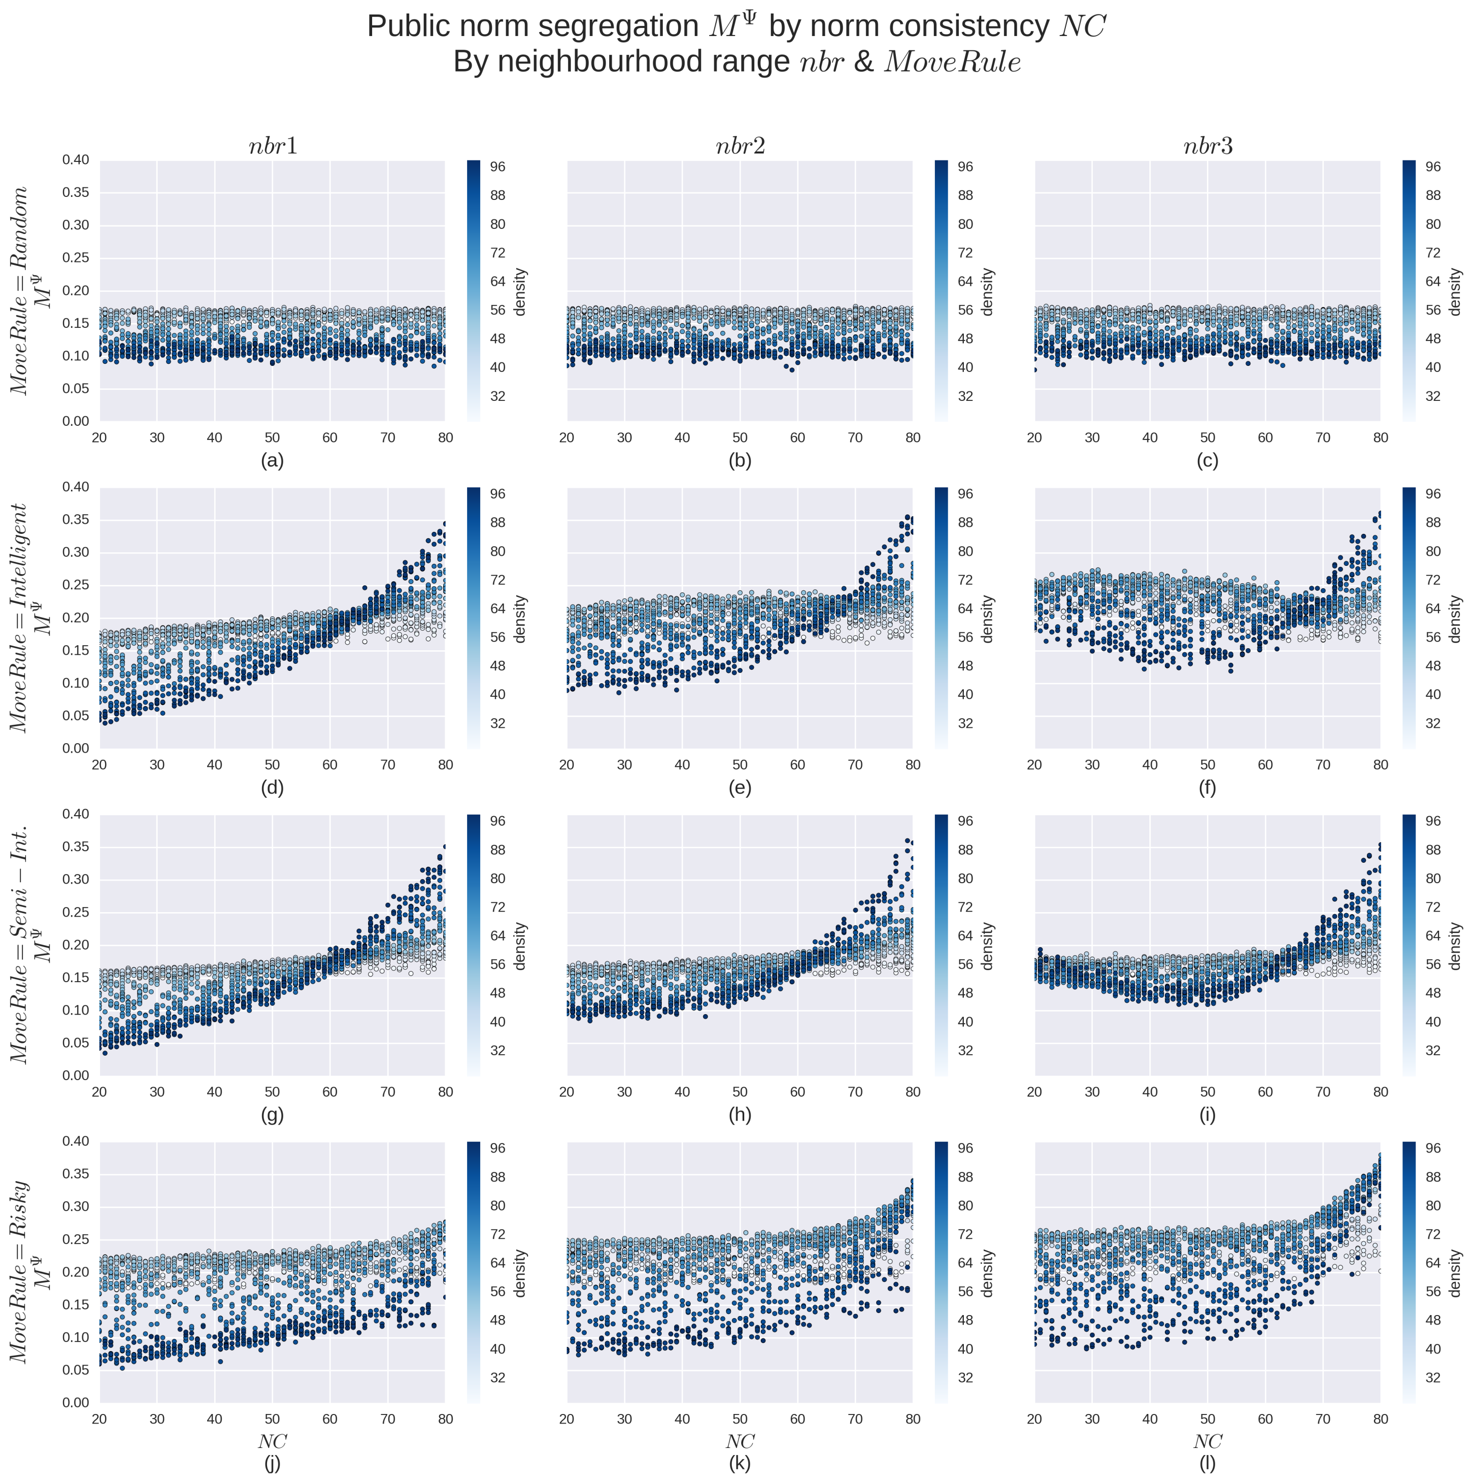

vate 'beneath the surface' pattern is more clustered. The scatterplots thus confirm the visual impression that was given by Figure 4.2 in Section 4.4.1. As truthfulness increases, the two measures of segregation measure increasingly the same thing, leading to the greater similarities between private and public norm segregation levels at high levels of NC.

4.6 Discussion

In this chapter I sought to determine what behaviours lead to a high level of PI, and what behaviours mitigate its occurrence. I have tested various reasoning patterns in order to distinguish the effects of reliance on external signals and the desire to belong to a group.

One major element concerns the combination of IC and M^{Ω}. Taking Figures 4.5 and 4.9 together, pluralistic ignorance is also prevalent under random movement, but clustering is much more common when agents do not move randomly. PI emerges automatically, possibly due to the strict definition outlined in this chapter. The combination with clustering when agents reason means that levels of fragility of assumed public opinion is higher: in clusters, agents reinforce their deception and misconception, and create closed spaces that leave little room for escape. In the migration society of a perceived public consensus around tolerance of minorities, this could mean that low-tolerance agents cluster together, but only some in the cluster signal their true preference- enough to attract deceivers, but too few to repel those who disagree with the public projection. This cluster is vulnerable to exposure of the false consensus: a norm entrepreneur may start to challenge the public perception, and there are many agents within the cluster that could then flip. Moving is costly, so any sudden shift in opinions will take time to break up again. This presents a vulnerability to populists that can target people who silence themselves. Populists can use their rejection of established institutions and authorities to present themselves as a norm entrepreneur that sends dissenting signals out, without the urge to silence themselves (Nome and Weidmann, 2013).

A second finding is that when agents reason whether to move, they are most effective in reducing PI and keeping costs low. I recall that the lattice in this model represents social space, not physical space- reinforced by the limited range of agent movement. They can 'jump' to empty tiles, but only those within the reach of their neighbourhood. To change friends is a costly endeavour. Thus, movement should be lower rather than higher. Figure 4.5 demonstrates that when agents decide whether or not do move, they will move far less on average than those agents who do so randomly. At peak movement levels, no more than 20% of all intelligent or semi-intelligent agents move. Semi-intelligent agents have been the most effective strategists in preventing pluralistic ignorance: they move little and will often decide not to move- but when they do, they will chose a random new location. By not trying to find norm-friends, they inadvertently prevent PI situations. These agents thus rely on signals only when determining their own state of comfort, but not when deciding where to go. This pattern persists for all but the highest levels of density. In very dense social space, it is impossible to avoid deceivers.

The model generates a range of stylistic facts that are reminiscent of recent debate on media consumption, 'filter bubbles' and reliance on social media (Groshek and Koc-Michalska, 2017). By avoiding to follow just like-minded people and by diversifying the social realm, pluralistic ignorance has less of a chance to occur: when people openly disagree, there is no false consensus. Reducing reliance on indirect signals (analogous to preventing a too-crowded situation) could involve reducing the consumption of media through others, such as social media or opinion pieces in newspapers.

The model in this chapter has used simple movement rules. Rather than switching private attitudes or bending to other kinds of pressure such as global opinions or strength of attitudes, agents solely seek to minimise the difference between their surrounding neighbours and themselves. The cost of movement through its restriction to the current neighbourhood emulates the high cost of switching- while switching does happen in the real world (but not in this model), the purpose of this design was to identify situations in which a sudden wave of flipping would have the most fertile ground. This is when agents are too crowded, many of which sending false signals, and a heavy reliance on signals of other agents. We already know that PI is emergent. I have shown that the combinations of relying on signals (or cues) and movement costs exposes groups to situations in which flipping could occur.

Pluralistic ignorance occurs not because people are actively striving towards a dominance of norm, but because every person tries to reduce any possible internal friction that might arise from conforming to a norm they do not privately agree with. These micro-behaviours lead to the macro pattern of pluralistic ignorance. The answer to how pluralistic ignorance emerges is thus: people bring it onto themselves. There is no need for any external pressure through coercion or legal norms, no active exchange between agents (that might result in peer pressure type settings) is required either.

The pattern of pluralistic ignorance is a value-neutral phenomenon which can occur as a catalyst for both positive and negative norm-change and preservation. In the case of student drinking and peer pressure, or rules that harm individuals, pluralistic ignorance is a hindrance to morally superior norms, if one assumes that individual happiness or opportunities are the goal. But it can also occur to bring change such as the changing norms about racist thought- white supremacy was once a widely accepted norm, but in most western countries today this has shifted. Racism still exists, but the attitudes underlying racism are no longer the accepted norm, even though they are contested (Inglehart et al., 2016). For such large changes to happen, it is more than plausible to imagine that there has to exist a period of pluralistic ignorance in which dissidents of the new norm find themselves in a position in which they are ostracised, and knowing what the new norm has become, will shift to identifying non-breaching of the new norm as an implicit compliance of the new norm, despite not knowing what the others' private attitude may be.

In the past, pluralistic ignorance could explain how high levels of racism could be sustained in the US, despite researchers' knowledge that groups of white people had changed their private attitudes (O'Gorman and Garry, 1976). The best 'cure' for prejudices remains direct contact with the out-group, as has been repeatedly demonstrated in the effect of exposure to black Americans through military enlistment (Fischer et al., 2016). However, it is not possible for all groups of people to be in direct contact with all other relevant groups about which they may harbour prejudices. The aim should perhaps be to view consensuses with healthy suspicion, and to identify proxies of dissent- for instance, general distrust towards foreigners by generally white, male and less educated citizens can be an indicator for a more general disagreement with cultural shifts in society (Inglehart et al., 2016).

Finding more stringent conditions for groups of pluralistic ignorance theory is important to distinguish the phenomenon from social acquiescence bias and other related concepts. I have proposed one such possible definition here, which hopefully can serve as a starting point for further discussion on this matter.

5 Conclusion

Social norms, tolerance, and social norms of tolerance have become increasingly important in a world which is increasingly diversifying. International migration shows no sign of abating, and existing societies that had previously built self-identity on ethnicity see their beliefs challenged by the growing diversity. The US is projected to turn into a majority-minority population by 2043 (Rios and Wynn, 2016). Europe, traditionally Christian, experiences continuing decline of Christian faith, coupled with increasing immigration from Muslim countries Joppke and Torpey (2013). Western society has seen its previous cultural shifts towards multiculturalism, pluralism and has started to undermine privileges that parts of the populations used to enjoy, sparking subsequent backlashes (Inglehart et al., 2016). The increased diversity in the US leads, contrary to previously suggestions, not to an increase in the Democratic vote share, but to a stronger partisan divide as marginalised groups and those who have experienced (or perceive) a loss of privilege are driven towards Conservative views (Craig and Richeson, 2014). In the smaller social realms, diversity is still dividing scholarship on when and where its effects are good or bad.

In the context of all of these changes and challenges (and opportunities), I have looked at several aspects of these changing dynamics with the intention to shed more light on the mechanisms that drive changes of tolerance, reactions to migration and sudden shifts in opinion about the norms that guide our understanding of tolerance and what we ought to believe. I have focused on the use of Agent-based models to explore potential explanations to further the existing empirical and theoretical debates in Political Science.

In this final section of the thesis, I will reiterate the main findings of each chapter in relation to the research questions and discuss the implications of the findings. This is accompanied by a review of agent-based modelling as a complementary method for studying migration. I conclude by reviewing the scope of this thesis and further work that can be carried out on the basis of my findings.

5.1 Main findings

One purpose of this thesis has been to apply agent-based modelling to the broad research area of immigration, more specifically, migrant-receiving societies. The topic of immigration spans many disciplines and is a vast subject area with many new developments and approaches. Recent migration research in Political Science has focused on the impact that diversity has on previously homogeneous or largely homogeneous societies. The results of these studies vary greatly, driven by different levels of analysis and operationalisation of social variables such as trust, prejudice, social capital or ethnic diversity. Agent-based modelling is a computational method based on the principles of complexity theory. It is designed to implement agents, an environment, and the rules governing the agents. On this micro-level, rules can be directly implemented, rather than inferred from aggregates. The interaction of agents can result in macro-level patterns that are emergent, rather than pre-determined. Using this approach, complex systems can be modelled without having to understand all of its properties. Human society is one such complex system, and immigration or changes in ethnic diversity of populations are changes induced into these complex systems. Using this approach, I have modelled societies that are on the receiving end of migration. Agent-based models of migration are numerous, but have thus far concentrated on either migrants themselves, or the flow of migration, not on the host society (Klabunde and Willekens, 2016a). I thus have aimed to bring agent-based modelling and recent research on population diversity and immigrant sentiment together to approach the topical questions of how migration impacts a society, and what role norm conformity can play in such a society.

5.1.1 How does migration affect host society and migrant community?

Two-thirds of the thesis have been devoted to the physical impact of migration and how societies of migrants and natives react and adjust both in the short-term and long-term. The research question is intentionally broad, to reflect the complexity-and agent-based approach that I have used, basing model parameters and rules on well-known stylistic facts, existing models and social theory.

The first section introduced migration to the Schelling model of ethnic segregation as the baseline model with which the research questions would be tackled. The Schelling model is well-established and robust, making it easier to compare my findings to other Schelling approaches, but also providing a legitimate basis for some of the model decisions. The ethnic segregation nature of the model made it an intuitive choice when testing the impact of (ethnic) diversity on a given population. The first version of the model in Chapter 2 represents a parsimonious effort to asses the effects that migrants have on the local 'native' population in the host country, and how migrants experience their newfound situation. I introduced immigration to the Schelling model: agents would no longer be spawned at the start of the simulation, but rather the model distinguishes between 'natives', initialised in the beginning, and 'migrants', agents that enter the grid at a later stage. The focus of this chapter was to address two particular concerns regarding migration: how, if at all, does the rate of change of migration impact the happiness and segregation behaviour of the overall population? And does the placement of migrants, i.e. the manner in which they arrive, change these happiness and segregation levels differently? These are frequently debated subjects in diversity and migration research (Laurence and Bentley, 2016). By providing a virtual repeated experiment, parameters can be taken to the extreme without ethical considerations that would be present in a real-life experiment.

Because this parsimonious model is not data-driven, any conclusions can be drawn as analogies or suggestive of mechanisms that drive macro-level patterns that can be picked up empirically.

The Schelling model is largely robust to the experimental treatments of migration, but exhibits short-term volatility and a sharp reduction in agent happiness when migration waves occur. Large, one-off influxes of migrants are difficult to absorb by the host society initially, and overall happiness and segregation levels are worse compared to more steady introduction of migrants. The short-term shocks in the model can be relevant for policy decisions. Electoral cycles for example are short, and decades-away potentially positive developments may not be as useful for politicians to use in an election campaign or to justify policy. It would thus be in the politicians' interest to introduce migration slowly and steadily, if an electoral cycle is part of the consideration. Even the introduction of diasporas must not necessarily prompt negative reactions- in the model, diaspora migrants were the happiest, and natives coped just as well with the existence of diasporas. This pattern is consistent with theories that assume that only when critical levels of diaspora size are reached, do natives 'turn' (Collier, 2013). The differences between rates of change of migration mirrors the empirical findings that have been made in the wake of Brexit: not the level of migration mattered in predicting a vote for leave, but rather the rate at which an area had been exposed to newcomers (Inglehart and Norris, 2017). That migrants are most happy in a diaspora is unsurprising- especially when only few migrants exist in the host society, actively targeting people of similar/shared backgrounds can be a relief when language barriers exist or initial settlement is otherwise difficult (Chong, 2000).

Important for the outcomes of the model were the levels of intolerance, F_2: as long as these agents preferred no more than 20 out of 24 neighbours to be of their own kind, overall happiness was much higher for both migrants and natives. Any values beyond 20 resulted in sharp dips in happiness for natives. The implication is that degrees of intolerance matter. From a policy perspective, this could mean that rather than providing a binary choice between yes-migration and no-migration, sceptical voters for whom migration is a concern, could be targeted to promote benefits of migration whilst acknowledging that the *perception* of a threat of migrants is very real for these voters. Most recent work on populism and immigration finds that a problem for pro-immigration parties is that their voters do not perceive migration as a problem, and for these parties migration is a non-issue, enabling anti-migration or migration-sceptic parties and politicians to frame the issue (Lambert et al., 2017). The mechanisms employed in this model suggest that even among high intolerance, small gains can come a long way.

The second version of this migration model in Chapter 3 has introduced adaptive tolerance behaviour, modelled after the contact theory, prominent in the field of intergroup relationships and out-group prejudice. This adaptation has aimed to demonstrate the ability of agent-based models to implement existing social theories and scrutinise their individual-level assumptions, but also to expand the parsimonious immigration model to include more realistic mechanisms without overwhelming the model with additional layers of attributes and interactions.

Adaptive behaviour is a crucial component of complex systems and in the deterministic environment of computer models, we can trace the causal mechanisms that produce the model results. Using the adaptive tolerance model, I have found a strong tendency towards bimodality of tolerance: over time, society of migrants and natives will split into very tolerant and very intolerant camps. Very small minorities are more likely to become tolerant, as they find too little room to cluster- but when a critical minority size is reached, intolerant clusters can be sustained. Just as integration of out-group members is an unstable state, so is moderate tolerance. Moderate tolerance is mostly observed short-term throughout the model when intolerant clusters turn tolerant or vice versa. Crucially, this behaviour is robust with regards to very slow rates of change of tolerance up to $\Delta_f = 0.001$. Even under these conditions, the system will tend towards bimodality. This outcome mirrors some of the debates in the political landscape of 2016/17 in Western countries such as the US, the UK, or other European states such as the Netherlands (see Goodwin and Heath (2016) on the UK, Inglehart and Norris (2017) on the US and Heath et al. (2016) for Europe more generally). The 'liberal divide' may be as inevitable as segregation itself, according to this model. If too few contact situations come about, isolated parts of society will tend towards intolerance. If one were to construe a tolerant, inclusive society, one would have to make sure that contact is frequent, but limited in scope. This avoids triggering threat responses and negative contact situations. Additionally, the spread of out-groups must be ensured so that isolated parts of society do not remain isolated for too long. This model shows that these mechanisms give rise to the divided society reminiscent of actual countries, but it cannot ascertain that there are not other possible explanations and mechanisms that can also achieve this.

Three key assumptions of the contact theory were modelled: that positive contact drives up tolerance, that contact is not frictionless, and that negative contact drives down tolerance (in the form of segregating behaviour). These rules were sufficient in generating a society with the parallels to real-world societies mentioned above: there were no additional variables required such as income, housing restrictions, type of labour, left-right positions or religious and ethnic background that would usually feature in questions around voting behaviour and anti-immigrant votes. The model mechanics suggest that the contact theory provides a good framework to understand how micro-level contact between groups can, on the aggregate level, lead to tolerant and intolerant pockets of society.

5.1.2 How does norm conformity interact with discontent of status quo?

The final third of the thesis moves on from purely physical impacts and norm of migration to the realms of social norms and engages in a theoretical exercise. Norms are oftentimes the reason why friction can occur between migrants and natives: both groups have an inherent set of social norms that eases communication and cooperation among groups, but inhibits it between different groups (Chong, 2000). The subject of interest of this research question concerns itself with the norms that exist around integration of migrants or inclusion and tolerance of minorities in general. I have looked at the phenomenon of pluralistic ignorance (PI), a situation that describes a collective mistake of a group to assume a status quo that is in fact a minority view. PI can explain how previously seemingly unchanging public attitudes can flip in a short period of time (Shamir and Shamir, 1997). The phenomenon is relevant to the political landscape of today, which has seen the issue of migration and integration of migrants rise sharply in this decade and that social norms of minority inclusion have potentially dampened outspoken criticism of migration (Blinder et al., 2013). The chapter has furthermore highlighted how agent-based models can be used to support a theoretical argument that I had laid out previously.

The model of emerging pluralistic ignorance in Chapter 4 highlights the ability of agent-based models to not just model, but advance theory. The wide range of loose definitions of pluralistic ignorance is a source of confusion for scholars that take interest in the subject. By incorporating PI in an ABM, I was forced to specify each element of the existing theory. This has resulted in a strict definition of PI which enables the ability to replicate the model and consider the effects of each of the assumptions that was made to devise this working definition. I have demonstrated that pluralistic ignorance can emerge unwittingly through the interplay of wanting to belong to a group of like-minded norm conformers and the inability of agents to determine what their neighbours *actually* think. This is a possible explanation for the rise of attitudes against migrants and populist parties that build on concerns surrounding migration (Blinder et al., 2013). It also throws up parallels to the model from Chapter 3. Clustering is once again an important driver of social behaviours that may not be intended. Clustering of deceiving agents is fragile: perceived norm consensus can be broken up quickly, but underlying attitude shifts remain slow-paced. For policy makers, this can signal that perceived consensus might not always be stable. If we suspect that inaccurate predictions of private preferences are a major culprit, it can open up opportunities such as amending survey questions to incorporate this inaccuracy(as has been done for other models- see O'Gorman

(1986)). If party-surveying prior to election campaigns included the possibility of voter deception with their own voters or members, target messages could be altered accordingly.

Public opinion shifts or changes in media narrative can appear sudden, whereas in fact they may have been no longer genuine and simply perpetuated through a false consensus effect. It thus takes a norm entrepreneur to shock the system of pluralistic ignorance into changing the narrative. Populists are such norm entrepreneurs: they reject the legitimacy of the status quo and accept the costs of publicly dissenting against prevailing norms. Since the late 20th century, these approaches have found fertile ground in many countries (Norris and Inglehart, 2018). It is thus useful to identify potential situations which could be misleading. A similar effect has taken place in the study of elections: the prediction of several recent elections has been increasingly difficult due to the mobilisation of previously inactive voters and generally a disconnect between self-reported likelihood of voting and actual turnout (Mellon and Prosser, 2015). If consensus of norms and views (such as post-modernism) can be viewed with the same healthy suspicion, false consensus situations may be detected.

5.2 The scope of this thesis and future work

The preceding section has summarised the main findings of this thesis, which has applied agent-based modelling and a complexity-theory perspective to the issue of migration and its impact on segregation, tolerance, and norms.

Whilst the explanations suggested by the models can inform policy makers, they cannot serve as proof that the suggested mechanism is in fact the only possible explanation. Thus, to advance policy proposals, the proposed mechanisms should first be validated by empirical work. Agent-based modelling can disentangle micro-macro connections and suggest re-directions for empirical research, but it is not designed to replace it. ABM is, due to its experimental nature, very useful for theory testing and discussion, as I have shown in Chapters 3 and 4. In non-empirical literature, agent-based models can provide an alternative method for proof of concept in contrast to equation-based models.

The initial Schelling model with immigration introduced in Chapter 2 was a parsimonious model basing the model design on stylistic facts, empirical findings and social theory, where available. A virtue and weakness of agent-based models is that they can be expanded to incorporate ever-more complex interactions, attitudes and agents. The more moving parts are introduced in an ABM, the harder it becomes to track the causal mechanisms resulting in observed macro-patterns- the reason why the model was reduced to as few moving parts as possible. Any addition should be justified. One such addition would be the introduction of a $FShare$ parameter that handles the majority of tolerant agents over intolerant agents. Because group size is a critical factor in determining population happiness and segregation clusters, it stands to reason that tipping points and critical values also exist for the distribution of tolerance among non-adaptive agents. The potential benefit of this addition could be to see if and when minorities of intolerance can no longer be sustained, or when majorities of tolerance become fragile. From the perspective of a policy maker, knowing that segregation levels and/or happiness levels can change drastically with just little differences in tolerance distribution, it can inform migration policy or political campaigning.

The findings of this thesis and the models that have been devised provide opportunities for further study. One future project that builds on this thesis will focus on the migration dynamics in greater detail. Based on the findings of Chapter 2, the migration treatment interacts strongly with population density and majority-minority ratios within the population. The suggested project considers migration on a continuous scale, checking whether there are thresholds in rate and size of migration at which collective behaviour changes in some ways. In Chapter 2 I have already highlighted one such threshold (see Section 2.11) and more work needs to be done. In addition to this further exploration of the parameter ranges of the model, I will consider scenarios of absorption of migrant clusters into the host society over time, as suggested by Collier (2013). This will also include different diaspora sizes and the probability with which an existing migrant group can exert influence on future migrants of the same group. Heterogeneous absorption and integration levels have been suggested as early as the 1940s (Penninx, 2006b, p.134) and have not been modelled as such. The goal of this research is to perhaps determine possible mechanisms of integration (or lack thereof) depending on the socio-demographic characteristics of both the immigration groups and host society communities. Furthermore, the model can serve as a baseline for further exploration of the impact of migration on inequality, which too has been a rising concern. For instance, the variance in happiness at each point in time in the model could serve as a proxy for inequality. Migrants and natives may show different variance levels, as has been indicated in the analysis of Figures 2.8 and 2.9 respectively.

Lastly, further work that includes agent-based modelling can incorporate empirical data as a means to validate the model. The parameter range can be cross-checked with existing data. For example, initial tolerance levels of agents could be adapted to tolerance levels of European citizens responding to a social survey. A growing field within the agent-based community is the use of GIS data. Agent-based models, instead of using 2-dimensional lattices with square tiles, use existing GIS data to shape the lattice with shapefiles that reflect the actual geographical boundaries of residential areas. Crooks (2010) have implemented such a version for the Schelling model, providing a useful basis for GIS-ABM work that builds on the abstract models of this thesis. A major difficulty facing GIS-based models is that the social variables are usually not gathered with high-resolution geographical information, so that combining social and geographical information requires extrapolating from larger aggregate findings. An alternative approach to using surveys is to rely on

GPS-enabled social media data such as Twitter feeds and networks, which have a growing relevance in studies of elections and voters.

The adaptive migration model could be adjusted to incorporate agent death or exit from the grid, modelling change of population composition over time. This can address theories put forward by Collier (2013) on the tendency of homogeneous, low-fertility societies relying on immigration (and thus, very likely, an introduction of ethnic diversity) to shift not just positions but composition. A commonly requested addition to the model is some form of distinction of income or class. For example, three types of agent could be introduced- low-, mid-, highly skilled or educated agents, that could reflect the divide between pro-migration cosmopolitans and migration-sceptic, provincial voters. There is evidence that skills in particular changes public perception of migrants greatly (Lambert et al., 2017), even trumping concerns about ethnic differences. Among natives, the "ethnic penalty" is reduced or eliminated when the assumed migrant is highly skilled (Lambert et al., 2017). An inclusion of skills or more broadly socio-economic differences would enable a better fit with existing economic models of migration and be beneficial to bridge a gap between social models and economic models of migration and its relevance to host societies.

In this thesis I have sought to demonstrate three useful ways in which agent-based modelling can further research in Political Science. It can provide experiments on models of parsimony, exploring basic mechanisms (Chapter 2). It can implement existing social theory directly, scrutinising its suggested mechanisms and predictions (Chapter 3), and lastly it can be used to support theoretical advancements in areas that have previously relied on equation-based arguments (Chapter 4). The virtue of agent-based models in Political Science is that it can complement areas in which empirical research has been slow or hindered by practical restrictions, although it should not be understood as a replacement empirical research.

The cultural backlash to societal changes in the West, the election of Trump, the vote for Brexit, the general surge of populists and the fears of terrorism, demographic shifts and concerns about globalisation are all significant events and developments that have marked the 21th century so far and will, at least in the near future, continue to present challenges for all elements of society. In this context, Political Science and the social sciences more generally can be crucial in providing explanations of how different causes and effects are linked, and how they function. In this context, I have provided potential explanations of such mechanisms which contribute to the growing body of knowledge of tolerance and norms in migrant societies.

References

Abid, R. Z., Manan, S. A., and Rahman, Z. A. A. A. (2017). A flood of Syrians has slowed to a trickle: The use of metaphors in the representation of Syrian refugees in the online media news reports of host and non-host countries. *Discourse and Communication*, 11(2):121–140.

Adam, H. and Moodley, K. (2014). *Imagined Liberation: Xenophobia, Citizenship and Identity in South Africa, Germany and Canada*. Temple University Press.

Allport, G. W. (1954). *The Nature of Prejudice*, volume 35. Addison-Wesley Publishing, 25th anniv edition.

Andersson, R., Brattbakk, I., and Vaattovaara, M. (2017). Natives' opinions on ethnic residential segregation and neighbourhood diversity in Helsinki, Oslo and Stockholm. *Housing Studies*, 32(4):491–516.

Andreouli, E., Greenland, K., and Howarth, C. (2016). I don't think racism is that bad any more': Exploring the end of racism' discourse among students in English schools. *European Journal of Social Psychology*, 46(2):171–184.

Arbaci, S. (2007). Ethnic segregation, housing systems and welfare regimes in Europe. *European Journal of Housing Policy*, 7(4):401–433.

Ariely, G. (2014). Does diversity Erode social cohesion? Conceptual and methodological issues. *Political Studies*, 62(3):573–595.

Arzheimer, K. (2009). Contextual factors and the extreme right vote in Western Europe, 1980-2002. *American Journal of Political Science*, 53(2):259–275.

Baganha, M. I., Doomernik, J., Fassmann, H., Gsir, S., Hofmann, M., Jandl, M., Kraler, A., Neske, M., and Reeger, U. (2006). *International Migration and Its Regulation*, pages 19–40. Amsterdam University Press.

Beheshti, R. and Sukthankar, G. (2014). A Normative Agent-based Model for Predicting Smoking Cessation Trends. *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)2*.

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

Bjerring, J. C., Hansen, J. U., and Pedersen, N. J. L. L. (2014). On the rationality of pluralistic ignorance. *Synthese*, 191(11):2445–2470.

Black, R., Biao, X., Collyer, M., Engbersen, G., Heering, L., and Markova, E. (2006). *Migration and Development: Causes and Consequences*, pages 41–64. Amsterdam University Press.

Blinder, S., Ford, R., and Ivarsflaten, E. (2013). The Better Angels of Our Nature: How the Antiprejudice Norm Affects Policy and Party Preferences in Great Britain and Germany. *American Journal of Political Science*, 57(4):841–857.

Branscombe, N. R., Schmitt, M. T., and Schiffhauer, K. (2007). Racial attitudes in response to thoughts of white privilege. *European Journal of Social Psychology*, 37(2):203–215.

Brennan, G., Eriksson, L., Goodin, R. E., and Southwood, N. (2013). *Explaining Norms*, volume 1. Oxford University Press.

Bruch, E. and Atwell, J. (2015). Agent-Based Models in Empirical Social Research. *Sociological Methods & Research*, 44(2):186–221.

Bryne, D. (1998). *Complexity Theory and the Social Sciences: An Introduction*, volume 39. Routledge.

Bushi, M. (2014). Rethinking Heterolocalism: The Case of Place-Making among Albanian-Americans. *Geography Honors Projects*.

Centola, D., Willer, R., and Macy, M. (2005). The Emperor's Dilemma: A Computational Model of SelfEnforcing Norms. *American Journal of Sociology*, 110(4):1009–1040.

Chong, D. (2000). *Rational lives : norms and values in politics and society.* University of Chicago Press.

Citrin, J. (2015). Are We All Now Multiculturalists, Assimilationists, Both, or Neither? In *Migration*, pages 138–160. Oxford University Press.

Clark, W. a. V. and Fossett, M. (2008). Understanding the social context of the Schelling segregation model. *Proceedings of the National Academy of Sciences of the United States of America*, 105(11):4109–4114.

Clough, E. (2007). Strategic voting under conditions of uncertainty: A re-evaluation of duverger's law. *British Journal of Political Science*, 37(2):313–332.

Collier, P. (2013). *Exodus : how migration is changing our world.* Penguin Books, London, first edition.

Collins, A., Petty, M., Vernon-Bido, D., and Sherfey, S. (2015). A call to arms: Standards for agent-based modeling and simulation. *JASSS*, 18(3):2–12.

Cortez, V., Medina, P., Goles, E., Zarama, R., and Rica, S. (2015). Attractors, statistics and fluctuations of the dynamics of the Schelling's model for social segregation. *European Physical Journal B*, 88(1).

Costa-Lopes, R. U. I., Dovidio, J. F., Pereira, C. R. C. R., and Jost, J. T. (2013). Social psychological perspectives on the legitimation of social inequality: Past, present and future. *European Journal of Social Psychology*, 43:229–237.

Craig, M. A. and Richeson, J. A. (2014). On the precipice of a "majority-minority" America: Perceived status threat from the racial demographic shift affects white Americans' political ideology. *Psychological Science*, 25(6):1189–1197.

Crawford, J. T., Brandt, M. J., Inbar, Y., Chambers, J. R., and Motyl, M. (2017). Social and economic ideologies differentially predict prejudice across the political spectrum, but social issues are most divisive. *Journal of Personality and Social Psychology*, 112(3):383–412.

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.

De França, D. X. and Monteiro, M. B. (2013). Social norms and the expression of prejudice: The development of aversive racism in childhood. *European Journal of Social Psychology*, 43(4):263–271.

Dovidio, J. and Gaertner, S. (2000). Aversive Racism and Selective Decisions: 1989-1999. *Psychological Science*, 11:315–319.

Dovidio, J. F. and Gaertner, S. L. (2010). Intergroup bias. *The Handbook of Social Psychology*, pages 1084–1121.

Drinkwater, S., Kauser, R., and Crawley, H. (2013). Regional Variations in Attitudes Towards Refugees : Regional Variations in Attitudes Towards Refugees : Evidence from Great Britain. *Centre for Research and Analysis of Migration*, 26(7647):45.

Elliott, G., Rothenberg, T. J., and Stock, J. H. (1996). Efficient Tests for an Autoregressive Unit Root. *Econometrica*, 64(4):813–836.

Epstein, B. (2011). Agent-Based Modeling and the Fallacies of Individualism. In Imbert, C. and Humphreys, P., editors, *Models, Simulations, and Representations*, pages 115–144. Routledge, New York.

Epstein, J. (2006). *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton University Press, Princeton.

Epstein, J. M. and Axtell, R. (1996). *Growing artificial societies : social science from the bottom up*. Brookings Institution Press, Massachusetts.

Esses, V. M., Jackson, L. M., Dovidio, J. F., and Hodson, G. (2008). Instrumental Relations Among Groups: Group Competition, Conflict, and Prejudice. In *On the Nature of Prejudice*, chapter 14, pages 225–243. Blackwell Publishing.

Fields, J. M. and Schuman, H. (1976). Public Beliefs About the Beliefs of the Public. *Public Opinion Quarterly*, 40(4):427 –448.

Fischer, M. J., Lundquist, J. H., and Vachon, T. E. (2016). Residential segregation: The mitigating effects of past military experience. *Social Science Research*, 60:61–73.

Foner, N. and Alba, R. (2008). Immigrant Religion in the U.S. and Western Europe: Bridge or Barrier to Inclusion? *International Migration Review*, 42(2):360–392.

Ford, R. and Goodwin, M. (2017). A Nation Divided. *Journal of Democracy*, 28(1):17–30.

Gilbert, N. (2007). *Agent-Based Models*. SAGE.

Gilbert, N. and Conte, R. (1995). *Artificial Societies: The Computer Simulation of Social Life*. UCL Press.

Goodwin, M. and Heath, O. (2016). The 2016 Referendum, Brexit and the Left Behind: An Aggregate-Level Analysis of the Result. *The Political Quarterly, forthcoming*, pages 1–16.

Groshek, J. and Koc-Michalska, K. (2017). Helping populism win? Social media use, filter bubbles, and support for populist presidential candidates in the 2016 US election campaign. *Information Communication and Society*, 20(9):1389–1407.

Gunther, A. C., Perloff, R. M., and Tsfati, Y. (2008). Public Opinion and the Third-Person Effect. In *The SAGE Handbook of Public Opinion Research*, pages

184–192. SAGE Publications Ltd, 1 Oliver's Yard, 55 City Road, London EC1Y 1SP United Kingdom.

Hainmueller, J. and Hopkins, D. J. (2014). Public Attitudes Toward Immigration. *Annual Review of Political Science*, 17(1):225–249.

Hall, S. (2013). Super-diverse street: a 'trans-ethnography' across migrant localities. *Ethnic and Racial Studies*, 00(00):1–14.

Hatna, E. and Benenson, I. (2012). The Schelling Model of Ethnic Residential Dynamics : Beyond the Integrated - Segregated Dichotomy of Patterns. *Journal of Artificial Societies and Social Simulation*, 15(2012):1–23.

Hatna, E. and Benenson, I. (2015a). Combining segregation and integration: Schelling model dynamics for heterogeneous population. *Journal of Artificial Societies and Social Simulation (JASSS)*, 18(4):1–22.

Hatna, E. and Benenson, I. (2015b). Combining segregation and integration: Schelling model dynamics for heterogeneous population. *Jasss*, 18(4):1–22.

Hatton, T. J. (2005). Explaining trends in UK immigration. *Journal of Population Economics*, 18(4):719–740.

Heath, A., Richards, L., and The Questionnaire Design Team (2016). Attitudes towards immigration and thier antecedents: topline results from Round 7 of the European Social Survey. Technical Report 7, European Social Survey European Research Infrastructure Consortium.

Helbing, D. (2010). Pluralistic Modeling of Complex Systems. *Science*, 76(9-10):1–20.

Helbing, D. (2012). Agent-Based Modeling. In *Social Self-Organization*, pages 25–70. Springer.

Inglehart, R. and Norris, P. (2017). Trump and the populist authoritarian parties: The silent revolution in reverse. *Perspectives on Politics*, 15(2):443454.

Inglehart, R. F., Norris, P., and School, H. K. (2016). Trump, Brexit, and the Rise of Populism: Economic Have-Nots and Cultural Backlash Faculty Research Working Paper Series.

Joppke, C. and Torpey, J. (2013). *Legal integration of Islam : a transatlantic comparison*. Harvard University Press.

Kahneman, D. (2011). *Thinking, fast and slow*. Penguin Books.

Katz, D. and Allport, F. H. (1931). *Students' Attitudes: A Report of the Syracuse University Reaction Study*. The Craftsman press, Syracruse.

Kaufmann, E. and Goodwin, M. (2016). The Diversity Wave: A meta-analysis of ethnic diversity, perceived threat and native white backlash. [note].

Kaufmann, E. and Harris, G. (2015). "White Flight" or Positive Contact? Local Diversity and Attitudes to Immigration in Britain. *Comparative Political Studies*, 48(12):1563–1590.

KhosraviNik, M. (2010). The representation of refugees, asylum seekers and immigrants in British newspapers. *Journal of Language and Politics*, 9(1):1–28.

Kitts, J. A. (2003). Egocentric Bias or Information Management? Selective Disclosure and the Social Roots of Norm Misperception. *Social Psychology Quarterly*, 66(3):222–237.

Klabunde, A. and Willekens, F. (2016a). Decision-Making in Agent-Based Models of Migration: State of the Art and Challenges. *European Journal of Population*, 32(1):73–97.

Klabunde, A. and Willekens, F. (2016b). Decision-Making in Agent-Based Models of Migration: State of the Art and Challenges. *European Journal of Population*, 32(1):73–97.

Krosnick, J. A. (1999). Survey research. *Annual review of psychology*, 50(1):537–567.

Kuran, T. (1995). *Private Truths, Public Lies. The Social Consequences of Preference Falsification.* Harvard University Press.

Labovitz, S. and Hagedorn, R. (1973). Measuring Social Norms. *The Pacific Sociological Review*, 16(3):283–303.

Lambert, J., Fitzgerald, R., Callens, M.-S., Dennison, J., Ford, R., and Hangartner, D. (2017). Attitudes towards immigration in Europe: myths and realities. Technical Report June, Migration Policy Group.

Laurence, J. and Bentley, L. (2016). Does ethnic diversity have a negative effect on attitudes towards the community? A longitudinal analysis of the causal claims within the ethnic diversity and social cohesion debate. *European Sociological Review*, 32(1):54–67.

Lerman, K., Yan, X., and Wu, X. Z. (2016). The "majority illusion" in social networks. *PLoS ONE*, 11(2):1–13.

Markaki, Y. and Longhi, S. (2012). What determines attitudes to immigration in european countries? an analysis at the regional level. CReAM Discussion Paper Series 1233, Centre for Research and Analysis of Migration (CReAM), Department of Economics, University College London.

Mclaren, L. M., Franklin, M., Hooghe, M., Morales, L., Letki, N., De Vries, C., Park, A., Evans, G., Van Der Meer, T., and Pardos-Prado, S. (2012). The Cultural Divide in Europe: Migration, Multiculturalism, and Political Trust. *the European Research Centre on Migration and Ethnic Relations.*

McPherson, M., Smith-lovin, L., and Cook, J. M. (2001). Homophily in Social Networks. *Annual Review of Sociology*, 27:415–444.

Mellon, J. and Prosser, C. (2015). Investigating the Great British Polling Miss: Evidence from the British Election Study. *Available at SSRN 2631165*.

Miller, D. (2000). *Citizenship and national identity*. Polity Press.

Miller, D. T. and McFarland, C. (1991). When social comparison goes awry: The case of pluralistic ignorance. In *Social comparison: Contemporary theory and research*, pages 287–313. Lawrence Erlbaum Associates, Inc.

Mitchell, M. (2009). *Complexitiy: A Guided Tour*. Oxford University Press, New York.

Moreland, R. L. (2010). Are Dyads Really Groups? *Small Group Research*, 41(2):251–267.

Moy, P. and Rinke, E. M. (2012). Attitudinal and behavioral consequences of published opinion polls. In *Opinion Polls and the Media*, pages 225–245. Palgrave Macmillan UK, London.

Nome, M. A. and Weidmann, N. B. (2013). Conflict diffusion via social identities : entrepreneurship and adaption. In Checkel, J. T., editor, *Transnational Dynamics of Civil War*, pages 173–204. Cambridge University Press, Cambridge.

Norris, P. and Inglehart, R. (2018). *Cultural Backlash: Trump, Brexit, And the Rise of Authoritarian Populism*. Cambridge University Press, New York.

Novotny, J. and Hasman, J. (2015). The Emergence of Regional Immigrant Concentrations in USA and Australia: A Spatial Relatedness Approach. *Plos One*, 10(5):e0126793.

O'Gorman, H. J. and Garry, S. L. (1976). Pluralistic ignorance - a replication and extension. *Public Opinion Quarterly*, 40(1976):449–458.

O'Gorman, J. H. (1986). The discovery of pluralistic ignorance: An ironic lesson. *Journal of the History of the Behavioral Sciences*, 22(4):333–347.

O'Sullivan, D. and Haklay, M. (2000). Agent-based models and individualism: Is the world agent-based? *Environment and Planning A*, 32(8):1409–1425.

Penninx, R. (2006a). *Conclusions and Directions for Research*, pages 305–318. Amsterdam University Press.

Penninx, R. (2006b). *Introduction*, pages 7–18. Amsterdam University Press.

Perryman, M. R., Davis, C. R., and Hull, S. J. (2017). Perceived community acceptance of same-sex marriage: Persuasive press, projection, and pluralistic ignorance. *International Journal of Public Opinion Research*.

Pettigrew, T. F., Christ, O., Wagner, U., and Stellmacher, J. (2007). Direct and indirect intergroup contact effects on prejudice: A normative interpretation. *International Journal of Intercultural Relations*, 31(4):411–425.

Pettigrew, T. F. and Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of personality and social psychology*, 90(5):751–783.

Pettigrew, T. F., Tropp, L. R., Wagner, U., and Christ, O. (2011). Recent advances in intergroup contact theory. *International Journal of Intercultural Relations*, 35(3):271–280.

Pfister, R., Schwarz, K. A., Janczyk, M., Dale, R., and Freeman, J. (2013). Good things peak in pairs: a note on the bimodality coefficient. *Frontiers in psychology*, 4:700.

Portes, A. and Vickstrom, E. (2015). Diversity, Social Capital, and Cohesion. In Dustman, C., editor, *Migration: Economic change, social challenge*, chapter 8, page 192. Oxford University Press, Oxford, first edition.

Prentice, D. a. and Miller, D. T. (1993). Pluralistic ignorance and alcohol use on campus: some consequences of misperceiving the social norm. *Journal of personality and social psychology*, 64(2):243–256.

Putnam, R. D. (2000). *Bowling alone : the collapse and revival of American community*. Simon & Schuster.

Putnam, R. D. (2007). E Pluribus Unum? Diversity and community in the twenty-first century. *Scandinavian Political Studies*, 30(2):137–174.

Ramos, M. R., Hewstone, M., Barreto, M., and Branscombe, N. R. (2016). The opportunities and challenges of diversity: Explaining its impact on individuals and groups. *European Journal of Social Psychology*, 46(7):793–806.

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics*, 21(4):25–34.

Rios, K. and Wynn, A. N. (2016). Engaging with diversity: Framing multiculturalism as a learning opportunity reduces prejudice among high White American identifiers. *European Journal of Social Psychology*, 46(7):854–865.

Rojas-Sosa, D. (2016). The denial of racism in Latina/o students' narratives about discrimination in the classroom. *Discourse and Society*, 27(1):69–94.

Sawyer, R. K. (2005). *Social emergence. Societies as complex systems*. Cambridge University Press.

Schanck, R. L. (1932). A study of a community and its groups and institutions conceived of as behaviors of individuals. *Psychological Monographs*, 43(2):i–133.

Schelling, T. C. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, 1:143–186.

Semyonov, M., Raijman, R., Yom-Tov, A., and Schmidt, P. (2004). Population size, perceive threat and exclusion: a multiple-indicators analysis of attitudes toward foreigners in Germany. *Social Science Research*, 33:681–701.

Shamir, J. and Shamir, M. (1997). Pluralistic Ignorance Across Issues and Over Time: Information Cues and Biases. *Public Opinion Quarterly*, 61(2):227.

Shin, J. K., Sayama, H., and Choi, S. R. (2014). A state equation for the Schelling's segregation model. *Complex & Intelligent Systems*, 2(1):35–43.

Simon, H. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1):99–118.

Singh, A., Vainchtein, D., and Weiss, H. (2009). Schelling's segregation model: Parameters, scaling, and aggregation. *Demographic Research*, 21:341–366.

Singh, A., Vainchtein, D., and Weiss, H. (2011). Limit sets for natural extensions of Schelling's segregation model. *Communications in Nonlinear Science and Numerical Simulation*, 16(7):2822–2831.

Skifter Andersen, H., Andersson, R., Wessel, T., and Vilkama, K. (2016). The impact of housing policies and housing markets on ethnic spatial segregation: comparing the capital cities of four Nordic welfare states. *International Journal of Housing Policy*, 16(1):1–30.

Smith, E. R. and Conrey, F. R. (2007). Agent-Based Modeling: A New Approach for Theory Building in Social Psychology. *Personality and Social Psychology Review*, 11(1):87–104.

Sniderman, P. and Hagendoorn, L. (2007). *When Ways of Life Collide: Multiculturalism and Its Discontents in the Netherlands.* Princeton University Press.

Søholt, S. and Lynnebakke, B. (2015). Do Immigrants' Preferences for Neighbourhood Qualities Contribute to Segregation? The Case of Oslo. *Journal of Ethnic and Migration Studies*, 41(14):2314–2335.

South, S. J., Pais, J., and Crowder, K. (2011). Metropolitan influences on migration into poor and nonpoor neighborhoods. *Social Science Research*, 40(3):950–964.

Spohr, D. (2017). Fake news and ideological polarization: Filter bubbles and selective exposure on social media. *Business Information Review*, 34(3):150–160.

Stephan, W. G., Boniecki, K. a., Ybarra, O., Bettencourt, a., Ervin, K. S., Jackson, L. a., McNatt, P. S., and Renfro, C. L. (2002). The Role of Threats in the Racial Attitudes of Blacks and Whites. *Personality and Social Psychology Bulletin*, 28(9):1242–1254.

Stephan, W. G. and Renfro, C. L. (2002). The Role of Threat in Intergroup Relations. In Mackie, D. M. and Smith, E. R., editors, *From Prejudice to Intergroup Emotions: Differentiated Reactions to Social Groups*, chapter 11, page 316. Psychology Press, 1 edition.

Stephan, W. G., Ybarra, O., and Morrison, K. R. (2009). Intergroup Threat Theory. In Nelson, T. D., editor, *Handbook of Prejudice, Stereotyping, and Discrimination*, chapter 3. Taylor & Francis Group, New York, first edition.

Stimson, J. A. (2004). *Tides of Consent: How Public Opinion Shapes American Politics*, volume 4. Cambridge University Press.

Sunstein, C. R. (2003). *Why societies need dissent.* Harvard University Press.

UKCensus (2011). Uk gridded population based on census 2011 and land cover map 2007.

Urselmans, L. (2016). A Schelling model with immigration dynamics. In *Artificial*

Life and Intelligent Agents (ALIA): Second International Symposium, page 13. Springer.

Urselmans, L. (2017a). Source-code for schelling model with adaptive tolerance. `https://github.com/lursel/ABM-adaptive-tolerance`. [Online; accessed 26/3/2017].

Urselmans, L. (2017b). Visualisation of influx mechanism for migrants arriving in a schelling model. `http://videos.lurselmans.me/\#InfluxAlgorithm`. [Online; accessed 30/7/2017].

Urselmans, L. and Phelps, S. (2018). A Schelling Model with Adaptive Tolerance. *PLoS ONE*, pages 1–36.

Van Assche, J., Asbrock, F., Roets, A., and Kauff, M. (2018). Positive Neighborhood Norms Buffer Ethnic Diversity Effects on Neighborhood Dissatisfaction, Perceived Neighborhood Disadvantage, and Moving Intentions. *Personality and Social Psychology Bulletin*, page 014616721774476.

Van Boven, L. (2000). Pluralistic Ignorance and Political Correctness: The Case of Affirmative Action. *Political Psychology*, 21(2):267–276.

van Dijk, T. A. (1992). Discourse and the Denial of Racism. *Discourse & Society*, 3(1):87–118.

van Prooijen, J. W., Krouwel, A. P., and Emmer, J. (2017). Ideological Responses to the EU Refugee Crisis: The Left, the Right, and the Extremes. *Social Psychological and Personality Science*, 9(2):143–150.

Wang, S.-W., Huang, C.-Y., and Sun, C.-T. (2013). Modeling self-perception agents in an opinion dynamics propagation society. *Simulation*, 90(3):238–248.

Wasserman, H. (2010). 'We're not like that': Denial of racism in the Afrikaans press in South Africa. *Communicatio*, 36(1):20–36.

Wilensky, U. (1998). NetLogo Flocking model.

Willer, R., Kuwabara, K., and Macy, M. W. (2009). The False Enforcement of Unpopular Norms. *American Journal of Sociology*, 115(2):451–490.

Williams, K. D. (2010). Dyads Can Be Groups (and Often Are). *Small Group Research*.

Zhang, Q. and Jager, W. (2011). Agent based modeling of population dynamics in municipalities: Migration in the Derbyshire & Nottinghamshire cases in the UK. *Evolution*, 6(November 2008).