

AGENT BASED MODELLING IN SOCIAL PSYCHOLOGY

by

GREGORY CARSLAW

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PSYCHOLOGY

School of Psychology
College of Life and Environmental Sciences
University of Birmingham
August 2012

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

Agent based modelling is a tool that has been successful in developing theories in a wide range of fields, but its application to social psychology is still in its infancy. This body of work applies the agent based modelling method to areas of social psychology including contact theory, group dynamics, altruistic behaviour and social identity theory. In each of these areas an agent based model is introduced that furthers the relevant theories and taken together these models demonstrate the effectiveness of some of the techniques outlined in existing research as well as producing a unique recommendation for the applications of agent based modelling in social psychology. In the fourth and fifth chapters three existing agent based models are extended in line with multiple identity theory and doing so produces novel results that improve upon the explanations of the original models. Therefore it is concluded that for agent based modelling in social psychology it is important to always consider the impact of multiple identities upon our modelling efforts rather than always simulating the minimum group identities necessary to test a hypothesis.

Acknowledgements

Dietmar Heinke and Julie Christian, who guided me through this process for five years.

My family, whose love and support saw me through some difficult years.

Everyone who was there when it mattered in the summer of 2011.

The ESRC, who funded most of this research.

Table of Contents

Chapter 1: Introduction	1
The Emergence of Computational Modelling in Psychology	1
Agent Based Models	4
Agent Based Models in Social Psychology	7
Principles for Applying Agent Based Modelling to Social Psychology	10
Theoretical Areas	13
Stigmatisation	14
Social Group Size	15
Social Exclusion	16
Social Identity Theory	17
Multiple Social Identities	18
Paradox of Altruism	19
Outlook	20
Chapter 2: An Agent-Based Simulation of Destigmatisation	21
Agent-based modelling approaches to social phenomena	22
The Social Psychological Framework for the Study	22
The Current Research	25
Destigmatisation Model (DSIM)	26
Study One: Basic Behaviour of DSIM	31
Method	31
Results	32
Discussion	38
Study Two: Levels of Stigmatisation and Group Perceptions	40
Method	40
Results	42
Discussion	44
Study Three: A Direct Manipulation of the Amount of Contact	45
Method	45
Results	46
Discussion	47
Study Four: Testing Allport's Moderators in DSIM	49
Method	49
Analysis	50
Results	52
General Destigmatisation Behaviour	52
Status	53
Authority	55
Cooperation	57
Common Goals	59
Interaction Analysis	61
General Discussion	65
Limitation and Outlook	69
Chapter 3: Modelling Fixed Social Group Sizes	71
Social Network Size	71

Social Psychology and Group Size	73
The Current Research	77
Method	78
Dynamic Groups Model (DGM)	78
Simulations	83
Data Analysis	84
Results	85
Chaotic Simulations	85
Example Simulation	86
Group Size	89
Coherency and Integration	92
Identification	94
Network Propagation	96
Discussion	100
Group Size Outcome	101
Identification Outcome	103
Network Propagation Outcome	103
Limitations and Outlook	105
Conclusions	106
Chapter 4: Extending Agent Based Models using Multiple Social Identities	108
Multiple Social Identities	108
The Current Research	109
Study One: Building on the Dynamic Groups Model	111
Method	112
Extended Dynamic Groups Model (E-DGM)	112
Outcomes	115
Results	115
Multiple Identities	116
Group Size	119
Identification	122
Proliferation	126
Discussion	128
Propegation	129
Multiple Identity Formation	129
Group Size	130
Study Two: A Classic Model of Racial Segregation	132
Existing Extension	132
Theoretical Background	133
Method	133
The Multiple Group Segregation Model (MGSM)	133
Initialisation and Data Treatment	135
Results	136
Replication	136
Additional Social Group	137
Balanced Multiple Identities	138
Limited Multiple Identities	141
Discussion	146

Reinterpretation of Schelling	146
Effect of Cross-Cutting Groups	147
General Discussion	148
Chapter 5: Extending Evolutionary Simulations with Multiple Societies	150
Evolutionary Computation	150
The Paradox of Altruism	152
The Present Research	155
Method	156
Baseline Simulation	156
Reputation Condition	162
Intersociety Competition Model (ICEM)	163
Results	165
General Pattern of Evolution	165
Important Emergent Behaviours	168
Replication	172
Intersociety Competition: No Reputation	173
Intersociety Competition: Friendly Forever	176
Intersociety Competition: Basic Reputation	179
Discussion	181
Chapter 6: Conclusions	184
General Discussion	184
Theoretical Results	184
The Contact Hypothesis	187
Social Group Size	190
Segregation	192
Altruism	195
Methodological Results	197
Theory Building	199
Sensitivity	201
Stable State Convergence	203
Multiple Groups	205
Limitations and Outlook	207
Final Conclusions	209
Appendix A: Survey Wording	211
Appendix B: Source Code	212
References	253

List of Illustrations

2.1 Average opinions of the two groups on each other over the course of the simulation.	33
2.2 Six exemplary time courses of the opinions of two majority agents, an "isolated majority" agent and a "contact majority" agent, and a minority agent.	35
2.3: Relationship between the minority size and the final stigmatisation of that minority.	37
2.4: Relationship between the number of minority agents and the opinion held of the minority in a typical simulation.	38
2.5 Different neighbourhood sizes	46
2.6: The effect of amount of intergroup contact on final stigmatisation score	47
2.7: Stigmatisation over time for different levels of common goals.	51
2.8a: Effects of Status on Overall Levels of Stigmatisation.	54
2.8b: Effects of Status on Destigmatisation Speed.	54
2.9a: Effects of Authority on Overall Levels of Stigmatisation.	56
2.9b: Effects of Authority of Destigmatisation speed.	56
2.10a: Effects of Cooperation on Overall Levels of Stigmatisation.	58
2.10b: Effects of Cooperation on Destigmatisation Speed.	58
2.11a: Effects of Common Interests on Overall Levels of Stigmatisation.	60
2.11b: Effects of Common Interests on Destigmatisation Speed.	60
3.1: Simulation state at cycles 0, 1, 50 and 100	87
3.2: Effect of 'similarity preference' and 'positive interactions' on group size.	89
3.3: Comparison of simulation progression for high positive interactions (80) and low positive interactions (20) at average similarity desire (50) at cycles 1, 50 and 200.	91
3.4: Effects of 'similarity preference' and 'positive interactions' on segregation.	93
3.5: Average Identification	95
3.6: Conditions under which the group moves beyond its initial area.	97
3.7: A 'positive interactions 80 similarity desire 80' simulation at 1, 25 and 50 cycles	98
3.8: A 'positive interactions' 20 'similarity desire' 20 simulation at 1, 50 and 200 cycles.	99
4.1: Agents with multiple identities.	117
4.2: Group Size outcomes for DGM full E-DGM and low similarity preference E-DGM	120
4.3: Identification Outcomes outcomes for DGM and E-DGM	123
4.4: Similarity Preference 50 Positive Interactions 20 simulation	125
4.5: Number of agents outside of the initial area for DGM and E-DGM	127
4.6: Replication of Schelling (1971)	137
4.7: Extension of Schelling with an additional social group	138
4.8: Average segregation over all agents at different levels of multiple identities.	140
4.9: Segregation results demonstrating interchangeability of the red, green and blue groups.	141
4.10: Multiple Ids 40 Preference 30 Cycle 32,715	143
4.11: The effects of agents with both red and green identities	144
4.12: Effects of red/green agents on blue segregation.	145
4.13: Effects of red/green agents on red and green segregation.	146
5.1: Sample population over time for 500 cycles	169

List of Tables

2.1: Survey Results	43
2.2: Levels of moderators in Study Four	50
2.3: Regression Analysis of Final Stigmatisation	61
2.4: Regression Analysis for Destigmatisation Speed	63
5.1: Frequency of different types of stealing rules in the final cycle of the simulation.	170
5.2: Results for altruism replication	172
5.3: Baseline results for inter-society competition by wins.	174
5.4: Baseline results for inter-society competition by losses.	174
5.5: Correlation matrix for inter-society competition without reputation	175
5.6: Correlation matrix for inter-society competition for "friendly forever"	177
5.7: Correlation matrix for inter-society competition for the basic reputation condition.	179
6.1 Theoretical contributions of the models.	186
6.2 Methodological contributions of the models.	198

List of Abbreviations

ABM

Agent based model. A model in which a large number of agents engage in relatively simple interactions with each other in the context of some environment. This thesis concerns the use of such models in social psychology, agent based models are introduced on page 4.

DGM

Dynamic Groups Model. A model introduced in Chapter 3 to explore the ways in which a group's size is formed by the behaviour of its members. The model is specified on page 78.

DSIM

Destigmatisation Simulation. A model introduced in Chapter 2 to expand upon the contact explanation for destigmatisation. The model is specified on page 26.

E-DGM

Extended Dynamic Groups model. An expansion of the DGM introduced in Chapter 4 to explore the effects of extending existing agent based models to account for findings in the multiple identification literature. The model is specified on page 112.

ICEM

Intersociety Competition Evolutionary Model. An expansion of Spronck and Berendsen (2009) introduced in Chapter 5 to explore the effects of extending existing agent based models to account for findings in the multiple identification literature. The model is specified on page 163.

MGSM

Multiple Group Segregation Model. An expansion of Schelling (1971) introduced in Chapter 4 to explore the effects of extending existing agent based models to account for findings in the multiple identification literature. The model is specified on page 133.

SIT

Social Identity Theory. Tajfel and Turners (1979) theory describing how individuals relate to social groups. The theory is introduced on page 16.

SFP

Self Fulfilling Prophecies. Used in the context of self-fulfilling prophecy theory (Jussim et al. 2000) which describes a method by which stigmatisation might be maintained. The theory is introduced on page 26.

UNHCR

The Office of the United Nations High Commissioner for Refugees. One of the organisations that supplied group size data for comparison to the survey reported in Chapter 2. This data is reported on page 43.

Social Agent Based Modelling

Chapter One: Introduction

Computational processing powers continue to improve dramatically, showing no sign of slowing down. Over the last ten years the processing speed and memory capacity of computers has increased by a factor of more than thirty. The ongoing development of computational resources has been mirrored by the rapid development of methods for furthering scientific inquiry using these tools. The nature of psychological investigations, particularly the difficulties in quantifying some types of finding, initially acted as a barrier to the uptake of these methods compared to the "harder" sciences. However the number of computational models published in psychological journals has risen in recent years and journals dedicated solely to this line of research have emerged.

This thesis concerns the application of computational methods to social psychology and looks to expand upon our knowledge from experimental evidence using these tools, as well as refining the tools for future applications. This chapter provides an overview of the ways in which psychologists use computational tools to develop and test their theories. The aim is to provide a clear backdrop before discussing, in detail, how the paradigm will be applied throughout the thesis and defining the domains within social sciences that it will be applied to.

The Emergence of Computational Modelling in Psychology

Computational techniques have been steadily developing and a variety of techniques are available to psychologists in this area, this section provides a brief overview of the

emergence of some of the key techniques. Early computational models extended the mathematical modelling paradigm. This approach is nothing new; the journal of mathematical psychology has been publishing papers concerning such models since its first issue in 1957. However, such computational tools enable much more efficient processing of such models. Recently mathematical models have been used to evaluate people's concepts of fairness (Gill & Stone, 2010), examine the impact of social network density on academic success (Rizzuto, LeDoux & Hatala, 2009) and a model of workplace altruism (Dur & Sol, 2010). These models continue to further our understanding of human behaviour, but do not differ in any fundamental way from the mathematical modelling approaches that have been common over the past century.

Parallel developments have also taken place in neurology. Mathematical models have been extended to model individual neurons within the brain and later modelling networks of these neurons (see Farley & Clarks, 1954). Increasingly complex neural network models have had applications in a variety of areas in psychology including, organizational identification (Lane Scott, 2007), language acquisition in children (Li, 2008), and person perception and social contact (Eiser, Stafford & Fazio, 2009). These models have contributed to the development of the field. However, they are constrained by our growing understanding of neurology and are not well suited to modelling the social psychology problems tackled in this thesis.

Evolutionary models play to the strengths of the modelling paradigm by taking advantage of the capacity to simulate thousands of simple interactions in a short space of time. In these models successful individuals pass on their attributes to the next generation, where unsuccessful ones do not, leading to each successive generation being better adapted to solve the problem presented by the model. While their primary application is in disciplines

such as biology and computer science, evolutionary models have contributed to the psychological literature as they have been used to explore the roots of observed human and animal behaviour in order to explain otherwise maladaptive behaviours. Recent examples include explanations of deception in aggressive competition (Szalai & Szamado 2009), group norm formation (Muller & Vercounter, 2010) and cooperative social learning (Segbroeck et al. 2010). The role of evolutionary psychology in general is contested (For a review see Badcock, 2012) but regardless of the ongoing debates surrounding the field, evolutionary models have provided valuable theoretical contributions.

Social simulations directly simulate societies in order to answer questions about how groups behave and interact. The earliest of these models is Schelling's (1971) model of social segregation, which is widely recognised by computational psychologists, and is frequently replicated and extended (Henry, Pralat and Zhang, 2011). The social simulations that followed can be broadly classified as those developed to predict social behaviour of a specific group in a specific context, often for an economic or political benefit and those developed to further scientific understanding of human behaviours. Both types of model provide useful information for computational psychologists.

The first type of social simulation has a variety of concrete benefits and applied social models seek to predict a specific behaviour with this goal in mind. Examples include predicting voter reaction to political speeches (Kim, Taber & Lodge, 2009), modelling criminal behaviour (Zhao & Costello, 2008) and developing effective emergency housing policy for predicted natural disasters (Sato, 2010). These models provide some limited benefits for theoretical psychology, but as the behaviours they uncover are not often generalisable to the wider population. Additionally, improving their effectiveness is one of the motivating factors in this field of research.

The second type of social simulation aims to identify general behaviours rather than the behaviour of a target population. For instance, it has long been known that in group discussions individuals opinions converge towards the groups norm (Sherif, 1937). Simulation studies of opinion convergence have been able to model this process in much greater detail. Opinion dynamics models have extended a simple convergence rule to consider where polarisation might occur (Hegselmann & Krause, 2002) or to account for the impact of individuals having different levels of influence (Yonta & Ndoundam, 2009). Axelrod's (1997) milestone simulation of opinion dynamics examined why it is, that despite the widely recognised tendency of individual's opinions to converge, we still enjoy a wide diversity of cultures and ideas and showed how multiple cultures can emerge from opinion convergence. Other areas targeted by theoretical simulations include social networks, in which researchers seek to understand the relationship between individuals and their social groups in the context of individual relationships. These models include White, McQuaker and Salehi-Abari's (2008) model of expert influence; Corten and Busken's (2010) model of the emergence of social norms and Gong's (2010) model of language transmission between generations. While all of the types of model discussed here have relevant lessons for psychologists, it is these models that directly seek to improve our theoretical knowledge of human behaviour that have the most to contribute and this is the goal of models presented in chapters two through five.

Agent Based Models

The theoretical models presented above use a variety of techniques, such as mathematical models (Yonta & Ndoundam, 2009), evolutionary models (Gong, 2010) and agent based models (Axelrod, 1997). These techniques are interrelated and a lot of them can

be understood in terms of each other. For instance any computable model is technically a mathematical model and could be expressed as a one line equation, though this would be highly impractical. As this thesis will focus upon the agent based modelling paradigm how each of the approaches presented in the previous section relate to agent based modelling will be presented below. The agent based modelling approach is selected as it has been identified as offering a number of important advantages to the study of social psychology (See 'agent based models in social psychology' below). Many of the studies mentioned in the previous section make use of agent based modelling and as the paradigm is important to this work it is helpful to revisit them in this context to highlight some of the strengths and weaknesses of agent based modelling as an approach.

An agent based model is a model that functions by having a large number of agents performing relatively simple behaviours in the context of some environment. The agents' behaviours influence each other and their environment, which in turn influences their behaviours. Commonly, agents are individuals and the environment is 'space' in which they interact, but these are not requirements. Agents can be any level of entity and an environment need not be spatial in nature. The goal of the model is for this cycle to reflect some observed phenomenon, so that the theories that drove the implementation of the agent's behaviour can be used as an explanation for the observed phenomenon. Agent based modelling overlaps with a number of other modelling techniques and domains (described above). Below I discuss the relationship agent based modelling has to mathematical modelling, neural network simulations, evolutionary simulations and social simulations to highlight the versatility of the approach, and some of the paradigms advantages that will be exploited in subsequent chapters.

The mathematical models discussed above generally did not use this method, though

there are models in which the equation used to represent an individual's state takes input from many other copies of itself (other agents) allowing for agent-based mathematical modelling. For example Snijders, Bunt and Steglich (2010) demonstrate how a multiagent mathematical model can reproduce the results of social network dynamics studies.

All neural network models are agent based models, which demonstrates an important principle of agent based modelling. In agent based modelling, it is key to remember that the definition of an agent is some smaller part of a more complex object. Often psychologists treat each agent as an individual, and look at the impact on a society or a species. But, it is just as valid to treat each neuron as an agent and the brain as the whole that is being studied. (Some biological processes parallel human behaviours and this similarity can be better used to understand both phenomena, as in the discussion of altruistic entities in chapter five.)

Evolutionary models that rely on genetic programming are exclusively agent based models as they rely on many individuals influenced by an environment created by other individuals. In other words, the act of selecting an agent as "the most fit agent" requires it be compared to other agents. So, in essence, it is situated in an environment defined by which other agents are present. Evolutionary simulations using this approach are widespread, from the numerous implementations and extensions of Lotka's (1920) classic predator-prey model to more recent evolutionary models such as Szalai and Szamados (2009) and Lammers, Warburton and Cribbs (2010).

Applied social models use a wide variety of methods to achieve their goals, agent based models have been used for tasks well suited to their strengths, such as the crime simulation by Furtado et. al. (2009). This model used agents representing police teams, criminals and the criminals targets in order to produce what they term a 'biologically inspired model of criminal activity'. Opinion dynamics models also tend to be agent-based models,

which each agent representing one opinion-holder. This is the case in the examples cited above (DeGroot, 1974, Hegselmann & Krause, 2002, Yonta & Ndoundam, 2009). Finally social networking simulations almost exclusively use the agent based modelling method, as studying a network by looking at its component individuals provides a powerful tool for examining the features of the social network. All of the studies mentioned in this paragraph use agent based modelling, either through a computerised simulation or a mathematical model. Features of social networking simulations are borrowed for the spatial models presented in chapters two through five, these are discussed in more detail in the appropriate chapters.

Agent based modelling has been widely applied to a lot of different areas of computational modelling within psychology. A particular strength of the method, when looking at aspects of social psychology, is that it allows the results of hundreds of interactions to be measured. Improvements in computational power have made it possible to examine many combinations of complex behaviours to arrive at the best explanations for the relationships between individual behaviour and large scale results. The higher level impact of psychological mechanisms observed in experimental studies are of interest to psychologists, but cannot practically be explored in survey or experimental settings.

ABMs in Social Psychology

While ABMs have been used extensively in other areas of biological and social sciences, application has been more limited within the field of social psychology. Using ABMs to drive optimisation algorithms has no real application in psychology as it is rare for

competing mathematical explanations in psychology to reach the level of complexity necessary to defy traditional mathematical approaches, largely due to the relative levels of precision obtainable in typical psychological experiments as opposed to, for instance, particle physics. Using ABMs to simulate particular social groups, as some of the models described in the previous section do, is equally unhelpful. This is because social psychologists are principally concerned with identifying general rules of human social behaviour, as opposed to predicting the behaviour of a specific group of people in a certain place and time. However the theory building approach to using agent based modelling does have a lot of potential in social psychology, as one of the perennial problems of the discipline is relating findings obtained from researchers operating under different sets of assumptions.

Smith and Conrey (2007) argue that ABMs could be fruitfully applied as a theory building tool in social psychology. They note that there are significant advantages to applying ABMs to this domain. Social psychologists often try to relate their research to macro (society) level issues, such as conflicts between social groups, socially learned behaviours and the challenges that individuals face due to belonging to a particular group. However, practical concerns mean that theories on these subjects must almost always be developed by gathering evidence from a restricted number of participants, and subsequently generalising to large group domains. Agent based modelling is perfectly suited to closing this gap, as it makes it possible to implement individual level behaviour and directly observe the implications of this on a large scale. It is also free from ethical and practical concerns that might limit other methods of study. For example it is possible to study the effects of disadvantaging a group without influencing well-being or other important aspects of function. Additionally, as with other types of modelling, it is powerful and enables for the testing of populations that would be difficult to access using traditional approaches.

The nature of specifying a simulation requires that every aspect of theory used to drive the simulation must be stated in an unambiguous way and any results found in a simulation are specified in a similarly precise way. This allows theory building agent based models able to produce very clear statements of the form: Under these theoretical assumptions the following properties emerge. Where the theoretical assumptions of a model are well evidenced in traditional psychology and the emergent properties have been observed in the wider world the approach yields theories that use existing psychological evidence to explain new phenomena. Alternatively if the emerging property has been observed as part of a different psychological theory it makes it possible to uncover relationships between existing theories.

Iannaccone and Makowsky (2007) provide a concrete example of the potential of agent based modelling to explain human behaviour. They observed that America can be geographically divided along religious lines, with highly religious areas in the south and less religious areas in the north. Existing explanations hinged on the idea that religion was a motivating factor when people were choosing where to live, however this contradicted survey data from people who had recently moved house, which showed that people considered religion to be relatively unimportant in making housing decisions. To examine this situation they created MARS (Multi-agent religious simulation) to look for other possible explanations. Iannaccone and Makowsky (2007) produced a MARS simulation in which this situation occurred. Religious regions persisted despite agents' moving randomly. MARS agents could do this, as they demonstrated a level of religious conformity to their neighbours. They conclude that the phenomenon that was observed (religious regionalism despite internal mobility) can be explained by an existing theory (conformity). While this could have been hypothesised without an ABM evidencing the point would have proven difficult as most

people are highly resistant to the idea that their religious choices might be largely based upon their neighbours and a study attempting to systematically manipulate an individual's religious beliefs would have encountered serious logistical and ethical problems.

Iannaccone & Makowsky (2007) provided a detailed example of the potential for agent based modelling in social psychology and over the course of this chapter other ABMs that have contributed to social psychology have been cited (for example, Muller & Vercouter, 2010, Schelling, 1971, Zhao & Costello, 2008, Bosse, Gerritsen & Treur, 2011, Axelrod, 1997 and Corten & Buskens, 2010). In this thesis I intend to expand on this body of work by developing new agent based models and expanding existing ones to further the use of this methodology in social psychology.

Principles for Applying Agent Based Modelling to Social Psychology

Designing agent-based models for the development of theories in social psychology requires the researcher to make a great number of decisions. Every aspect of the agents and their world must be specified in advance, but often existing research does not provide clear support for one implementation over another. To offset this conflict, there are some useful principles that modellers can use to guide their modelling efforts.

The most important of these is the idea of "minimal modelling". Any ABM intended for theory building should model the minimum number of factors required to model the target phenomenon. There are three reasons for this. Firstly, modelling a lot of factors is very time consuming and can lead to conflicting findings as factors need to be checked against each other. Secondly, there is a danger of over-fitting the available data on the target phenomena.

Having a lot of variables with complex mathematical relationships makes it almost certain that there will be some arrangement of these variables that matches the data, producing spurious results. Finally, and most importantly, in order to draw theoretical conclusions from ABMs, the researcher needs to relate the assumptions that underlie the model to the observed phenomena. If the model has a large number of assumptions and factors built in then drawing clear conclusions becomes impossible (see Smith & Conrey, 2007). This is distinct from the more general principle observed in Occam's razor, that explanations requiring the fewest assumptions should be preferred, in that it refers to the number of parameters explicitly modelled rather than the number of assumptions underlying the model. The relationship between the two means that observing the minimal modelling principle often results in models with a small set of assumptions, which is as desirable in agent based modelling as it is in science in general.

However, while observing minimal modelling, equally it is important to consider what impact factors not included in the model might have. While an additional factor that might lead to some additional finding, those findings will be less clear theoretical contributions, due to the increase in assumptions required by the model. Yet if considering an additional factor would change the existing findings and the additional factor is likely to exist in the population being simulated, then it is better to include it. An example of this can be seen in Gotts and Polhill (2010) who demonstrated that, for the FEARLUS model, the results changed qualitatively with the size of the environment, a factor often fixed at an arbitrary level by modellers. Thus, they demonstrated a need to vary an additional parameter in examining this model. The number of possible extensions that could be made to any given model are very large, to the point that it would take astronomical amounts of time to test them all, however a good model should at least test (even if these tests go unreported) some of the most obvious

extensions to their model.

Another important simulation concept is that of replication. Being able to replicate results is a core part of any science and with simulations direct replication is trivial, as the program already exists and can be run again, producing the same results. However, this is not always sufficient as demonstrated in Radax and Rengs (2010). They attempt to replicate a prisoners dilemma ABM that was not fully specified by its original authors. By identifying all likely implementations of the unspecified elements of the model they produced several similar models from the same set of theoretical assumptions, identifying the one that produced results most similar to the original study as a means to derive the method used. This process has two implications for modellers, firstly as Radax and Rengs (2010) correctly observed, modellers should seek to describe their models in enough detail that a different programmer might implement the model and obtain the results. The second is that the findings from a particular model may be a consequence of the specific implementation, rather than a consequence of the assumptions themselves. Therefore, modellers should always be cautious in interpreting their results and researchers seeking to replicate agent based models should consider alternative implementations of the same assumptions to determine if the original conclusions hold up to scrutiny.

Finally, the use of statistics in agent-based modelling requires a level of caution. Traditional statistical approaches often focus on determining whether two groups behave in significantly different ways. In an ABM two groups being examined have often been implemented differently by the programmer, so it is trivial to say that there are measurable differences between them. Additionally replication can give false certainty, as it is easy to repeat a simulation a large number of times and get consistent results. However repeating a simulation only tests how reliably the simulations assumptions produce a particular result, it

does not test the validity of the assumptions underlying the simulation or how well these assumptions were implemented (Chapter five discusses issues of replicating and extending simulations in more detail). Thus, in agent based modelling simulations a measurable statistical effect is necessary, but not sufficient, to draw conclusions from an ABM. Any statistical finding should be supported by clearly identifying the chain of events that lead to the relationship between the implementation of the model and the results it produces, so that the results can be properly interpreted on a theoretical level.

Theoretical Areas

The next four chapters of this thesis introduce agent based models concerned with different areas of social psychology. The purpose of this section is to provide an overview of these domains. The caveat, however, is that the focus of the thesis is upon developing the application of the ABM method for use in social psychology, rather than on the specific areas within social and personality psychology, thus the literature review is more focused in nature.

The models presented in the thesis are as follows: Chapter two introduces the destigmatisation simulation (DSIM), which simulates the destigmatisation process for stigmatised minorities groups. This examines the impact of Allport's (1954) contact moderators and demonstrates that they do not always facilitate each other. Chapter three develops the dynamic groups model (DGM), which serves to investigate the ways in which a social group's size might change in response to ingroup rejection, in order to relate findings on social network size (e.g. Hill and Dunbar 2002) to social group size. Next, chapter four extends two existing models, Schelling's (1971) model of social exclusion and the dynamic

group processes model (from chapter three). These extensions use the social identity and multiple social identity literatures to extend existing agent based models to take better advantage of the findings in social psychology. Finally, in chapter five, a similar extension is used to develop an evolutionary model to better address the paradox of altruism. The main theoretical areas relating to these models are outlined below and examined in greater detail in the relevant chapters.

Stigmatisation

In the context of group dynamics, stigmatisation describes a situation in which one group is the subject of negative attributions from another more powerful group. Stigmatisation has been shown to have consequences for members of the stigmatised group, including decreased mental and physical health, lower academic achievement and reduced access to jobs (See Major & O'Brien 2005 for a review of the effects of stigma). As the effects of stigma are almost universally agreed to be negative, a great deal of research has focused on looking for ways to reduce stigmatisation.

Chapter two introduces a model of a stigmatised group, and examines the ways in which that group might become destigmatised over time. The pattern of destigmatisation observed in the model closely mirrors a line of existing research started by Allport (1954). Allport proposed the contact hypothesis, which holds that contact between a nonstigmatised and stigmatised group would reduce stigmatisation under certain 'optimal conditions'. This theory has been developed to include a number of other factors and is still widely studied (For a meta-analysis see Pettigrew & Tropp, 2006). The model presented in chapter two expands

contact theory by demonstrating how the contact effect can emerge from implementations of other theories such as the self fulfilling prophecy account of stigmatisation maintenance (Jussim et al. 2000). I also test the effects of the optimal conditions in this context.

Social Group Size

A social network is a web of individuals, connected to each other directly through their personal relationships or indirectly through intermediaries. An individual's social network is defined by the set of people that they know. The definition of exactly what counts as a relationship for the sake of a social network has been a matter for debate, the size of an individual's social network might vary wildly depending on the method used to identify social links, for example Pool and Kochen (1978) tested several different methods producing estimates of average network size that varied from 1,100 to 4,250. Hill and Dunbar (2002) hypothesised that social network size was linked to neocortex size, based on evidence from primates, and using Christmas card lists as a measure of social network size found this to be true in humans. This famously produced “Dunbars Number” an estimate of the maximum social network size obtainable by the average human, approximately 150 individuals.

Besides feeding into the growing social networking literature, this finding is of interest to psychologists as Dunbars number has been popularised as applying to social groups. As well as being portrayed in the media this way (e.g. Allen, 2004) Dunbar directly references situations in which group sizes of 150 occur naturally (Dunbar, 2008). However this area has received little research attention and the factors that might cause social groups to be limited in this way have not been examined. Clearly, if this is the case, a social group does not have an

individual neocortex size. In chapter two a model of social group dynamics is presented that addresses these questions.

Social Exclusion

Social exclusion is distinct from stigmatisation in that it is concerned with individuals withdrawing social contact, and associated resources, rather than directly engaging in hostile behaviour. Social exclusion can be thought of as one aspect of stigmatisation, which describes a broader pattern (Abrams, Christian & Gordon, 2007). Social exclusion is of particular importance to agent based modelling, because initial modelling efforts focused on the emergence of social segregation in physical spaces (Schelling, 1971). The key finding of this model is that if two groups of agents had very mild preferences for being around their own group, the two groups would become completely segregated. This was used as an explanation for racial segregation in America at the time (which remains an issue today, Pettigrew, 2008).

This model of segregation was critical to the early development of agent based modelling as a field in social sciences and continues to inspire new research. For instance Henry, Pralat and Zhang (2011) apply this model in a social network context to show that homophily still emerges in most network topologies. In chapter four a model is introduced to develop the Schelling (1971) line of research further by investigating how the model's findings are influenced by multiple identity research (e.g. Crisp, Ensari, Hewstone & Miller, 2002, see multiple social identities below).

Social Identity Theory

Social identity theory (Tajfel & Turner, 1979) provides a framework for discussion on human social behaviours. The theory holds that individuals benefit from possessing a number of social identities. Social identity theory has been highly influential in the development of social psychology and has been extended throughout its lifetime. Turner (1987) expands it into self categorisation theory, which covers a broader set of phenomenon and postulates that that individuals position themselves with respect to social groups by comparing themselves with others. Hogg and Abrams (1998) and later Abrams and Hogg (2000) examine the developments in SIT over two decades, noting how effectively it has been applied to a wide variety of areas and suggesting future directions for the theory. Jetten and Postmes (2006) emphasises the role of the relationship between the group identity and the individual, moving the theory away from its roots as a theory of intergroup conflict. The theory that individuals have varying levels of identities with social groups has a lot of implications in agent-based models. Models of binary group membership are common and, as some of the models presented in this thesis will demonstrate, models that allow for varying levels of identity produce novel findings.

Research into the nature of individuals interactions with their social groups can be applied to a variety of agent based modelling endeavours. Agent-based models often concern the way in which individuals from social groups behave. These can either be specific to a single social group (e.g. The Italian groups modelled by Cracoli, Cuffaro & Nijkamps, 2009) or general in application (e.g. The abstract group relations modelled by Henry, Pralat & Zhang 2011). However, these studies do not take advantage of a rich literature describing human social behaviours, nor do they seek to account for the underlying psychological processes

which could be motivating the phenomena they seek to model. In fact, the level of detail available through existing research on social theories far exceeds our modelling capacities. Thus to implement these psychological models in the context of an agent-based model it is best to use a simplified theory and introduce other areas of the psychological literature only as required by the specific hypothesis under examination.

Multiple Social Identities

Social identity theory explicitly describes individuals as belonging to multiple social groups and researchers have examined how individuals approach having identities with contradictory norms. Brewer, Ho, Lee and Miller (1987) asked participants to make judgements upon individuals with two salient group memberships. The targets shared one identity with the participant but not the other. They found that participants still showed ingroup bias towards these individuals, but to a lesser extent than they did to targets who shared all group memberships. Crisp, Ensari, Hewstone & Miller (2002) identify several ways in which participants combined information from two groups to form judgements. In the absence of an experimental manipulation participants most commonly used the additive method, treating other individuals more favourably the more groups that they had in common. Chapter four takes this research and uses it to extend existing social models to account for individuals who hold multiple social identities, resulting in modifications to the original findings of the model.

Paradox of Altruism

The 'paradox of altruism' is a long contested problem in evolutionary biology, which is of interest to evolutionary psychologists. The challenge presented by the paradox of altruism is to explain how it is that altruism is so often observed in human and animal species, despite its apparent lack of value as a survival trait. In this context altruistic behaviour is taken to be any behaviour that decreases the fitness (chance to have successful offspring) of the active participant in order to increase the fitness of another individual. This is a biological definition, but also describes a many of the behaviours encapsulated by philosophical and psychological definitions.

Several explanations for the paradox of altruism have been proposed. Early explanations include reciprocal altruism (Trivers, 1971) and kin selection (Hamilton, 1964), however more recently multilevel selection (Wilson & Sober, 1994) has been advanced. This builds upon group selection theories to suggest that altruistic behaviour can survive by contributing to a higher level entity (e.g. a society) that is more capable of surviving than the sum of its parts. The relative roles of these competing explanations are still under investigation and the paradox of altruism is of interest to contemporary researchers (e.g. Warneken & Tomasello, 2009).

Multi-agent models have contributed significantly to the development of these theories, as for reasons stated in the evolutionary simulations section above they are well suited to modelling this domain. A biological model by Nahum, Harding and Kerr (2011) demonstrates that multiple bacteria cultures can develop altruistic behaviour as part of a systems level effect. Spronck and Berendsen (2009) examine the emergence of altruism in humans in terms of some of the theories described above. The fifth chapter of this thesis

expands upon this model by drawing upon social identity theory to produce a society level explanation consistent with multilevel selection theory.

Outlook

The areas identified above are all suitable for development through agent based modelling. They are each sufficiently developed to generate appropriate assumptions for the creation of an agent based model, but have sufficient unanswered questions that the model is able to be informative. The following chapters will each present a model based on the literature in one of these areas, taking existing findings forwards by asking “What are the consequences of these findings?”. Where the consequences of implementing a particular finding are novel ABMs are able to suggest fruitful avenues for future research. However the most useful results from ABMs resemble existing theories that were not used to drive the assumptions of the model, as these results are used to relate theoretical concepts to each other, forming a wider theoretical framework.

Chapter Two: An Agent-based Simulation of Destigmatisation

(DSIM): Contact theory and Self-fulfilling prophecy

Stigmatisation has long been recognized as a problem, and the continuing consequences of its negative effects on mental health (Harrell 2000), physical health (Clark et al. 1999), academic achievement (Steele 1997) and other areas such as social mobility and access to housing, education and employment are also well-established (See Major and O'Brien 2005 for a review of the negative effects of stigma). It is no surprise to find that one of the questions guiding much of the social psychology and social sciences literature on the topic has been to identify ways in which we can reduce intergroup bias and stigmatisation (Abrams, Christian & Gordon 2007; Wu et al. 2007; Finkelstein, Lapshin & Wasserman 2008). However the methodologies used in existing accounts do not allow researchers to simultaneously examine the individual level and group level processes that drive stigmatisation. A method that allows the interactions between individuals who belong to social groups to be related to the society level interactions between those groups would enable further development of this field. This chapter uses such an approach to advance our understanding of processes underlying destigmatisation by introducing a new theoretical approach in shape of an agent-based model (ABM).

Agent-based Model approaches to social phenomena¹

In an ABM a large number of “agents” interact with each other within an environment according to a set of simple rules. These models can examine how macro-level (social) phenomena emerge from micro-level (individual) interactions. Such models have been used to explore a variety of domains such as social segregation (Schelling, 1971), dating and mate selection (Kalik and Hamilton, 1986), opinion formation (Hegselmann and Krause, 2002), and religious choice (Iannaccone and Makowsky, 2007). These models generated novel explanations for these social phenomena. These explanations would have been difficult to unearth with traditional research methods. For example, Schelling (1971) demonstrated that total social segregation can emerge from very minor individual-level preferences in the social make-up of neighbourhoods. Iannaccone and Makowsky (2007) showed how religious preferences could emerge and be maintained, even in populations that frequently gain and lose members irrespective of religious orientation. In this paper we employ an ABM to the understanding of the reduction of social stigmatisation focusing on the interactions between individuals’ characteristics and how they influence the process of destigmatisation. This is termed the destigmatisation model (DSIM).

The Social Psychological Framework for the Study

In his seminal intergroup ‘contact hypothesis’ Allport (1954) suggests that contact between members of in-groups/mainstream groups and out-groups/stigmatised groups leads to

¹ This chapter has been submitted for publication, so there is some minor repetition of points recently made in the previous chapter.

a reduction in intergroup bias, a precursor to destigmatisation. According to Allport (1954) this occurs under what he considers to be “optimal conditions”, which should moderate this destigmatisation process. These optimal conditions are: *equal status between the majority and minority groups; intergroup cooperation; common goals/shared interests between the groups; and the support of authorities, law, or custom.*

Allport (1954) states that *equal status* is a necessary condition for contact to reduce bias. Later research, however, demonstrates that equal status is not necessary for intergroup contact to reduce prejudice, but increases the effectiveness of contact in prejudice reduction. For example, Jackman and Crane (1986) found that white Americans whom had regular contact with black Americans of high socio-economic status held more positive attitudes towards the black Americans than white Americans who had contact with those of a lower socio-economic status.

The "optimal condition" *cooperation* requires that intergroup contact occur in the context of a cooperative activity. Desforges et. al. (1997) examined the conditions under which personal prejudice reductions achieved by cooperation would best generalise to attitude change, using a confederate to take the role of a member of a group the participants had a pre-existing prejudice against (assessed in a pilot study). In this study cooperative interactions lead to positive attitude change. More recently, Oortwijn, Boekaerts, Vedder and Fortuin (2008) demonstrated that cooperative contact is effective in reducing intergroup conflict among children using a minimal group paradigm.

The recognition of *common goals* also has a positive effect upon the capacity of intergroup contact to reduce prejudice between members of differing social groups. It can be difficult to distinguish between cooperation and common goals, as most manipulations that supply a common goal also involve some form of cooperation. However, Chu and Griffey

(1985) studied players on interracial sports teams and found that members of teams requiring no direct cooperation (such as fencing teams) had no significant racial attitude differences to those requiring a cooperative effort (such as football teams).

Finally, Allport (1954) stated that intergroup contact would be most effective at reducing prejudice when the contact between the groups received support from *authorities* recognized by both groups. Parker (1968) examined interracial attitudes among the congregation of a mixed race church, and found five factors that had lead to good interracial attitudes in the church. One of these was the 'exceptional leadership' displayed by the minister.

Many of these factors were drawn together in Pettigrew and Tropp's (2006) meta-analysis. They systematically tested 713 independent samples from 515 intergroup contact studies. They confirmed the central hypothesis of the contact theory - that intergroup contacts decrease intergroup bias in a broad range of social settings - however, they also found that the optimal conditions improved destigmatisation, but were not a necessary precondition for it to occur. Moreover, their results suggested that *institutional support* had the strongest effect of the four conditions. However, they argued that this effect only makes sense if the status of social groups is similar. In contrast, if groups compete with each other (e.g. for institutional support), this support is more likely to have detrimental effect. From that, Pettigrew and Tropp (2006) concluded that the four conditions need to function together in order to have a facilitatory effect.

The Current Research

The destigmatisation model (DSIM) is an agent-based model developed in this chapter to examine the effects of intergroup contact. While there are many factors that influence the effectiveness of contact, DSIM is principally concerned with Allport's classic "optimal conditions". These conditions have repeatedly been shown to be important, most of the developments of the contact hypothesis in the previous chapter continue to examine them.

The remainder of the chapter is organised as follows: The next section introduces DSIM, including the implementation of Allport's conditions. The first set of simulations (Study One) demonstrate that DSIM implements a destigmatisation process. This process generalises across different group sizes (i.e. the minority group always becomes less stigmatised, irrespective of its size.) Interestingly, the simulations show that stigmatisation decreases with increasing group size. This prediction is empirically confirmed in Study Two, further validating the model. The reason for this prediction by DSIM stems from the fact that with increasing minority group size, the quantity of interactions with the majority group also increases. Hence, the contact effect emerges from the interactions between agents in DSIM before any of the modules of DSIM that explicitly implement the contact hypothesis are included. In the Study Three, the amount of contact is manipulated more directly by increasing the size of the neighbourhood in which agents interact with each other. This enables the analysis of the effects of more extreme levels of contact. At high levels of contact, the destigmatising effect of additional contact decreases and at some point additional contact no longer increases destigmatisation. Finally, Study Four examines the influence of Allport's (1954) four "optimal conditions" in DSIM, both in terms of the magnitude of their individual contributions to the destigmatisation process, and in terms of how they affect the

speed of destigmatisation.

Destigmatisation Model (DSIM)

The model consists of two groups of agents, a majority social group and minority social group. Each agent is placed on a spatial grid and is randomly allocated to either the mainstream unstigmatised majority group, or to the stigmatised minority group. Each agent possesses two numerical attributes, with a value ranging from -1 to 1. One value represents the agents' "opinion" of the unstigmatised majority group, while the other value represents the agents' "opinion" of the stigmatised minority group. The agents' opinion scores express their attitudes towards the respective group ranging from, "I really dislike this group" (-1) to "I really like this group" (1). In order to simulate individual interactions, agents interact with a subset of other agents (the method for selecting this subset varies between studies, see below). Agents interact in pairs and these social interactions can either be positive or negative. Importantly, however, both agents perceive a given interaction in the same way. Thus, up to this point, the set-up of the model follows a minimal modelling principle (e.g. only two "social groups" of agents and only single values for opinions about social groups).

Now, to realise how the interactions influence the agents' opinions. DSIM draws on the self-fulfilling prophecy theory of social interactions (see, Jussim et al. 2000; for a review of the supporting data). The self-fulfilling prophecy consists of three steps (e.g., Jussim, 1986). The steps are as follows: First an individual ("perceiver") develops expectations about others ("targets"); second, perceivers' expectations influence the way in which they treat the target individuals; third, targets react to this treatment with behaviour that confirms their

expectation. Note this separation into perceivers and targets is arbitrary, as in a didactic interaction both parties are perceivers and targets at the same time.

The self-fulfilling prophecy is implemented in DSIM in the following way: First, each agent generates an expectation of the quality of the interaction, which is equal to the opinion an agent holds of the other agent's social group. Second, the quality of the interaction (common to both agents) is influenced by this expectation about the interaction. However, in order to take into account the reciprocal nature of the interaction, the quality of the interaction also has to be affected by the other agent. This is implemented by averaging the opinions that the two agents have of each other:

$$\text{Quality of interaction} = (\text{opinion1} + \text{opinion2}) / 2 + \text{random}(-1, 1)$$

Random(-1,1) is a random number between -1 and 1 and takes into account that the quality of interactions, in reality, is influenced by numerous factors that we simply modelled as random noise. *Opinion1* is the opinion that the first agent holds of the second agent and *Opinion2* is the opinion that the second agent holds of the first agent. Finally, to complete the cycle of the self-fulfilling prophecy model, the opinion of the involved agents is changed through this simple rule:

If quality of interaction > opinion then increase opinion

If quality of interaction < opinion then decrease opinion

In other words, if the interaction turns out to be more positive than expected, the agent

improves its opinion of its partner's social group; however, if it turns out more negative than expected then its opinion of that social group decreases. (Opinions of the interaction generalise to the group.) Hence, the opinions move closer to the quality of interaction implementing the self-fulfilling prophecy. The change of opinion in all simulations was set to 0.0001. In pilot studies, this was found to be a reliable value, but one that also balanced issues of a convergence into a stable state of the model and reasonable length of simulation time.

Other aspects of DSIM vary between studies. In Study 1 agents interact only with other agents that are adjacent to them on the grid and the number of agents in the minority is varied. This is the simplest implementation of DSIM and the same overall result could be more efficiently obtained through traditional mathematical modelling, but such a model could not be so easily expanded in the subsequent studies and would not demonstrate the impact of spatial relations as successfully. In Studies 3 and 4 the distance at which agents will interact is varied and the minority size is fixed. These variations are discussed in detail in the relevant sections. Study 4 also introduces substantial changes to the model as it explicitly models the impact of Allport's optimal conditions. The conditions are: equal status between groups, common goals, intergroup cooperation and the support of authorities, law, or custom. The realisations of these conditions are expressed here as they are easier to follow in the context of the description given above, but note that they are not used until Study 4.

Status. The quality of interactions is determined using the average of the agents' opinions of each other as indicated above. Unequal status of group members is implemented by weighting the two opinions differently. When a member of the majority group interacts with a member of the stigmatised group the quality of interaction equation is modified as follows:

$$Quality\ of\ interaction = (opinion1 * (1 - status) + opinion2 * status) / 2 + random(-1,1)$$

Opinion1 represents the opinion of the majority group agent's opinion in the interaction. Here, *opinion2* is the opinion of the minority group agent in the interaction. The *status* parameter determines how equal the two groups are in status. It ranges from 0 (i.e., the minority agent's opinion has no effect on the interaction) to 0.5 (i.e., the minority agent's opinion has as much effect on the interaction as the majority agent's opinion does).

Common goals. To implement the 'common goals' parameter, the agents are arbitrarily assigned one of two abstract goals. Each agent is assigned to either goal "A" or goal "B". Agents retain the same goal throughout the simulation. The recognition of the goal is randomised, with the probability of recognition as a new parameter "common goals". In a given cycle, if agents recognise that they hold a common goal, those agents interact an additional time during that cycle (as an indication of liking for that other agent, because they share a 'common goal'). Agents do not remember one another's goals from cycle to cycle, so this test for recognition is made each cycle regardless of the outcome of the test in previous cycles.

Intergroup cooperation. The possibility of intergroup cooperation was implemented by designating some interactions to be 'cooperation attempts'. In these interactions both agents are attempting to cooperate, representing individuals risking some of their capital on a cooperative endeavour that could benefit both agents. In a 'cooperation attempt', more is at stake, so any change of opinion, whether positive or negative, is doubled. In order to vary the level of cooperation in a given simulation a new parameter describes the chance that each

interaction would be a 'cooperation attempt'.

Institutional support. To implement the notion of “institutional support” for intergroup interactions, some agents were designated as “key authorities”. Because institutional support is only effective when both social groups recognize the expertise of the authority, these authority agents belong to both the stigmatised and non-stigmatised groups, and therefore can function in both capacities. An important consequence of this, however, is that an authority agent never has an intergroup interaction itself. Rather, the sanctioning of intergroup interactions by these authority agents increasing the rate of intergroup interactions between minority and majority agents in their vicinity. That is, any intergroup interaction that occurs within three grid spaces of one or more authority agents will occur an additional time each cycle. The level of institutional support that occurred in a given simulation was varied by altering the number of authority agents placed at the start of the simulation.

Study One: Basic Behaviour of DSIM

The purpose of this study is to show that intergroup contact between agents that are members of stigmatised groups and agents that are members of non-stigmatised groups leads to destigmatisation. This will confirm that DSIM models the contact effect (Allport 1954). For DSIM to be successful, the simulations need to show a decrease in the stigmatisation of minority agents when there are larger minority sizes, as larger minority sizes indirectly increase intergroup contact.

Method

In Study 1 the model consisted of 22 x 22 grid, with a total of 484 agents divided into two groups. Simulations used one of four different sizes of minority groups. In a given simulation the minority group had: 30, 60, 90, or 120 agents. Agents interacted with their eight immediate neighbours. Simulations were run for 5,000 cycles; and each simulation was repeated ten times.

At the beginning of each simulation opinions were initialised to ensure that the minority was "stigmatised" and that the majority was "non-stigmatised". This was achieved by starting every agent's opinion of the minority group at -0.9. Every agent's opinion of the majority group was started at 0.9.

Results

Analysing the results for the model begins with understanding the models behaviour for a single parameter setting. The following results are for the simulations with a minority group size of 60, which were fairly typical of the models behaviour. Figure 2.1 shows the time course of the average opinions of the agents for one simulation. The majorities opinion of the minority rises from -0.9 to about -0.6 over the course of the simulation. In other words minority agents in DSIM became less stigmatised over the course of the simulation, making DSIM a model of destigmatisation. To examine whether this effect is statistically reliable over all simulations, the mean starting and final stigmatisation from each of the 40² simulations were compared by means of a t-test. The t-test revealed significant destigmatisation effect, $t(39)=13.83$, $p < 0.01$.

² Four group sizes were tested and each was replicated ten times.




Figure 2.1: Average opinions of the two groups on each other over the course of the simulation.

To understand why this happens with the groups as a whole, it is important to understand what happens to individual agents. Figure 2.2 presents the time courses of the opinions of two majority agents, an "isolated majority" agent and a "contact majority" agent, and a minority agent. The "contact majority" agent was adjacent to the minority agent. The "isolated majority" agent had no contact with the minority group. As shown in figures 2.2a and figure 2.2b, the "isolated majority" agent did not change its opinions of the minority group at all. It also did not change its opinions of the majority in any substantial way, while

its opinions fluctuated slightly they remained around the 0.9 mark at which they were initialised. Similarly the “contact majority” agent did not substantially change its opinions of the majority groups. The minority agent did not change its opinions of the minority group at all, as it had no other minority group to interact with. Finally and importantly, the “contact majority” agent and the minority agent both change their opinions of each other’s groups, converging towards each other’s opinions. These changes in opinion are what drive the results observed at the simulation level.


The reason for these opinion changes is made clear by examining the "Quality of Interaction" rule. If two agents that are interacting have the same opinion of each other then this rule simplifies to:

$$\textit{Quality of Interaction} = \textit{Shared Opinion} + \textit{random}(-1,1)$$


This means that the quality of interaction has an equal probability of being greater than or less than the existing opinion. Thus, while the opinion will change most cycles, changes are equally likely to be positive or negative, so over many cycles these random changes will cancel each other out leading to largely unchanged opinions. In contrast, for agents with different opinions of each other, the rule turns to:

$$\textit{Quality of Interaction} = \textit{Average Opinion} + \textit{random}(-1,1)$$


In this case the agent with the 'above average' opinion is more likely to decrease their opinion and the agent with the 'below average' opinion is more likely to increase it. Over many interactions the agent’s opinions will converge towards each other until they are equal and the stable state described above is reached.



a) Opinions of the majority



b) Opinions of the minority



c) Reciprocal Opinions of the Minority and Majority

Figure 2.2: Six exemplary time courses of the opinions of two majority agents, an "isolated majority" agent and a "contact majority" agent, and a minority agent. The "contact majority" agent (label: Contact) is adjacent to the minority agent. The "isolated majority" agent (label: Isolated) has no contact with the minority group (see main text for further discussions).

However from this explanation it is not completely obvious why the contact majority and minority agents converged towards a point closer to the majorities' negative opinion rather than the exact average of their opinions (See figure 2.2c). The explanation can be refined by accounting for the different sizes of the minority and majority groups. If a single minority agent is outnumbered by eight adjacent majority agents, then each cycle the minority agent interacts with the majority eight times (once with each majority agent) but each majority agent interacts with the minority only once. Thus the minority agent decreases its opinion eight times more quickly than the majority agents therefore the opinions converge more towards the negative side.

The temporal pattern of the individual opinions was observed in the overall simulation (see figure 2.1). However, unlike the individual level results, in which the contact majority and minority eventually form the same opinions of each other, the opinions of the majority and minority groups as a whole do not fully converge. This is due to the influence of the "isolated majority" agents who do not have contact with the minority therefore never change their initial negative opinions. This effect demonstrates that in DSIM intergroup contact is a necessary prerequisite to stigmatisation reduction. Hence DSIM simulates a crucial aspect of Allport's intergroup contact theory.

The impact of having sufficient minority agents to interact with a larger portion of the majority becomes more apparent upon comparing the results for the simulations at different minority sizes. Figure 2.3 shows how the average majority group's opinion of the minority group depends on the size of the minority group. The figure reveals that for all group sizes the minority group is less stigmatised than it was at the beginning of that simulation. It also shows that larger groups undergo more destigmatisation than smaller groups.




Figure 2.3: Relationship between the minority size and the final stigmatisation of that minority.

This pattern emerges due to increasing the number of minority agents in contact with the majority. This is confirmed by demonstrating the relationship between a majority agent's opinion of the minority and the number of minority agents adjacent to it (see figure 2.4).




Figure 2.4: Relationship between the number of minority agents and the opinion held of the minority in a typical simulation.

Discussion

Destigmatisation occurs in DSIM because initially majority agents 'expect' in-group interactions to be more negative than they are, which leads them to increase their opinions. The majority agent's expectations are flawed as they make a prediction without knowledge of the minority agent's opinion of them, which is initially higher than their opinion of the minority agent. This process can only take place for majority agents that experience intergroup interactions. The reverse process occurs for the minority agents, whose

expectations are initially too high. When the minority agent's opinions equal the majority agents opinions the process stops and the simulation becomes stable. As such if a majority agent has contact with a greater number of minority agents, it will undergo a larger opinion change, as multiple minority agents can undergo more interactions before their opinions are brought down to the majority agents level.

The simulation results confirm that the behaviour of the model is consistent with Allport's contact hypothesis, the full implications of which will be picked up in the general discussion. Only agents from the majority group that have contact with the minority group improve their opinions of the minority group and the more majority agents have contact with minority agents (as varied by group size) the smaller is the stigma at the end of the destigmatisation process. Note that the simulations demonstrate that for destigmatisation to occur in DSIM, it is essential that the minority group agents start out with a good opinion of the majority agents. If the minority group agents reciprocate the negative opinion of the majority group agents destigmatisation does not occur in DSIM. DSIM is a model of destigmatisation, and so is initialised such that the opinions held of the minority are lower than any other opinions. This assumption is made as opinions of a stigmatised group are, by definition, very low and because a group can become stigmatised without forming negative opinions of the majority. Stigmatised groups react to stigmatisation in a variety of ways other than forming negative opinions of the stigmatising group. For example women who do not perceive themselves as stigmatised (Lakhani, 2008) black reporters redefining stigma in the context of their professional identity (Slay and Smith, 2011) or homosexual couples strengthening their relationship in response to stigmatisation (Frost, 2011).

Study Two: Level of stigmatisation and Group Perceptions

The purpose of this study is to validate the previous chapters finding that there is a relationship between stigmatisation and population size. This will be achieved by conducting a survey to determine the extent to which different social groups are stigmatised so that this can be compared to population size data. Study One predicts that a negative relationship will be observed between the perceived stigma of a group and its size.

Method

Participants

Eighty participants were recruited through the University of Birmingham's official portal (www.bham.ac.uk portal). From the user base of this portal it is estimated that there were 36 male and 44 female participants with a mean age of 20.8 years.

Survey Measure

Using Sidanius and Pratto's (1999) stigma hierarchy work as a framework, participants were asked to rank order a list of 23 social groups, from high to low, in terms of how stigmatised they perceived each of the groups to be. Participants were instructed to assign each group a number between 1 and 23, using each number only once. In this context, a

lower number indicated a less stigmatised group, while a higher number indicated a greater degree of perceived stigma.

Survey Administration

The online advertisement posted on the portal contained a link to an online survey. The participants were told they would be participating in a study exploring social attitudes; and they were asked to complete the task as described above (see Appendices A for exact wording). All participants were informed that their results would be confidential, and would in no way impact their future relationship with the university. Upon completion they were thanked and debriefed.

Group Size Measure – Calculating Real World Population Sizes

The population parameters associated with of the 23 social groups were drawn from existing sources. The most common source used was the 2001 UK census, which supplied data for 14 of the groups. The other sources were government statistics (5 groups), relevant charities (3 groups) and a United Nations source, the UNHCR (1 group). The percentage of UK residents that belong to each of the groups was derived from this data, using the census to supply a total UK population size, and then later used for analysis.

Results

The results were analysed by calculating a stigmatisation rank for each social group, which was the average of the ranks assigned to it across individuals. In addition, each social group had a population size, expressed as the percentage, representing the percentage that this group forms of the UK population. (See Table 2.1)

The hypothesis that stigmatisation and group size were related was tested by examining the correlation between these variables. The results show an inverse relationship between group size and stigma, $r(22) = -.451$, $p < .05$, such that the more stigmatised the group the smaller the group as a total of the UK's population. Given that some data was drawn from advocacy groups, the analysis was repeated using only the UK Census data, because this is known to be the most accurate population data. (This reduces the analysis to 14 groups.) Again, there is a significant negative correlation between perceived stigma and group size, $r(13) = -.604$, $p < .05$. This supports the finding of the previous study, the more opportunities that one has for contact with stigmatised group (based on the larger number of people in the groups), the less stigma appears to be attached to those social groups.

Table 2.1: Survey Results

Group	Average stigmatisation Rank	Population Size	Population Size Source
People with cancer	4.04	3.3	Cancer Research UK
Victims of crime	5	26	www.statistics.gov.uk
Women	5.03	51.32	Census
Epileptic people	5.37	0.5	Epilepsy Research UK
People over 60 years old	8.23	20.76	Census
Jewish people	8.57	0.52	Census
Chinese people	8.75	0.45	Census
People of mixed race	8.96	1.31	Census
People registered as disabled	10.64	17.93	Census
Black Caribbean people	11.54	1.14	Census
Poor people	12.22	9.2	www.statistics.gov.uk
Indian people	12.26	2.09	Census
Gay people	12.55	6	Stonewall
Black African people	12.99	0.97	Census
Bangladeshi people	13.16	0.56	Census
People registered as unemployed	13.44	2.42	Census
Pakistani people	15.87	1.44	Census
Arabic people	16.04	0.37	Census
Refugees	17.2	0.5	UNHCR
People serving a jail sentence	17.72	0.14	www.statistics.gov.uk
Homeless people	18.21	0.18	www.statistics.gov.uk
Muslim people	18.25	3.1	Census
Immigrants	19.44	10	www.statistics.gov.uk

Discussion

This study was conducted to test whether the relationship between level of stigma and the associated size of the population found in Study One occurs in the general population. As indicated, smaller the populations were linked to greater levels of perceived stigma. While perceived stigma may differ from the actual level of stigma a group experiences, the two are strongly related, so the results indicate that there is a relationship between stigmatisation and group size. DSIM provides an explanation for why this effect occurs. Small populations within the more general UK population will have fewer chances for intergroup interactions with members of the mainstream, therefore keeping them more isolated, and thereby more stigmatised. This confirms the predictions that DSIM makes in study one and so validates the model. As such, further examination of the causal nature of the relationship is justified.

Study Three: A direct manipulation of the amount of contact

In order to study contact effects it is necessary to manipulate the amount of contact in DSIM. In Study One the increase of the minority group size lead to an increase of contact the majority group agents had with minority group agents. However this method conflates group size and the amount of contact. In order to being able to vary the amount of contact without changing group size, a new method is developed to manipulate contact. In Study One each agent interacted with its eight neighbours. In the present study the size of this neighbourhood will change. The aim of this study is to test whether this alternative implementation of contact maintains the relationship between contact and stigmatisation predicted by the contact hypothesis.

Method

Simulations were executed as in Study One, however in this study agents did not always interact with their eight immediate neighbours. Instead in a given simulation each agent was able to interact with 0, 2, 8, 16 or 24 of the agents in proximity to them in a fixed pattern (see Figure 2.5). The size of the minority group was fixed at 120 agents.




Figure 2.5 Different neighbourhood sizes (Interaction area highlighted)

Results

Testing whether a contact effect emerged required a comparison between the neighbourhood size in a simulation and the final opinion that the majority held about the minority in that simulation. Increased neighbourhood size lead to an increased opinion of the minority, though diminishing returns are evident (see Figure 2.6)




Figure 2.6: The effect of amount of intergroup contact on final stigmatisation score

Discussion

Generally speaking, the results indicated that increasing the size of neighbourhoods led to a decrease in the stigmatisation of the minority group agents. However, it is important to note that as neighbourhood size increases, this effect diminishes and ultimately saturates. This saturation can be understood in terms of the effect observed in Study One: A majority agent in contact with multiple minority agents forms a higher opinion of the minority, but as the number of minority agents in contact increased the effect diminishes (see figure 2.4). In Study One a cluster of eight majority agents surrounding one minority agent will converge to the same opinion. Moreover if one of those agents is adjacent to a second minority agent then

all of the involved agents will converge to the same opinion. The same is true of any majority agents that interact with the second minority agents, even if the neighbourhood size is too small for them to directly interact with any of the other ten agents. If the agents alternated between majority and minority every agent in the simulation would have the same intergroup opinions by the end of the simulation. As the agents are randomly placed, this is highly unlikely. Normally several clusters of agents whose opinions have converged will emerge. However as the neighbourhood size increases fewer clusters emerge, as it takes increasingly dense concentrations of one type of agent to prevent agents from two clusters interacting and causing the opinions of their clusters to converge. At the highest levels of neighbourhood size most agents are (indirectly) in contact with each other, so subsequent increases have little effect.

The simulation results demonstrate that the new manipulation of the amount of contact successfully mimics Allport's (1954) contact hypothesis. However and interestingly the simulations also predict that there should be a limitation in the influence of contact. The simulations suggest that at high level of contact additional contact does not reduce stigma. This has implications for Allports (1954) optimal conditions that are explored in the next study, the discussion of the saturation of the contact effect will be picked up in the general discussion.

Study Four: Testing Allport's four Moderators in DSIM

The DSIM studies so far focus on manipulating the amount of contact in the simulations. This is only sufficient to study simple contact effects. The extensions to DSIM, described previously, allow it to be expanded to test Allport's (1954) optimal conditions. This study will test whether the implementations of the moderators successfully mimic the empirical findings (see introduction). This test is conducted in two stages: The first stage establishes whether each individual moderator on its own improves the contact effect. (With contact operationalised by the neighbourhood size.) The second stage examines the interplay between the moderators in DSIM. This last set of simulations constitutes genuine predictions as empirical studies have not previously covered these interactions.

Method

The impact of the individual moderators was examined in a two-way factorial design. Each moderator had four levels (absent, low, medium and high, see Table 2.2). The second factor, neighbourhood size, was included with five levels (0, 2, 8, 16, and 32). The size of the minority was fixed at 120 agents. For each cell the simulations were repeated ten times, requiring a total of 800 simulations. This design allowed us to examine how DSIM behaves across different levels of the moderators. The simulations showed that all moderators across all level had an effect in a similar direction. Therefore only the highest and lowest level of each moderator was used in the interaction analysis. This reduction in levels also simplified the analysis of the results without losing generality.

Table 2.2: Levels of moderators in Study Four

Status	Cooperation	Authority Agents	Common Interests
0	0	0	0
10	10	10	10
30	50	50	50
50	90	90	90

Analysis

When destigmatisation occurs in the model it follows a pattern of initially rapid destigmatisation which tapers off until the agents opinions are no longer changing over time (see figure 2.7 for an example). These results are analysed by fitting a curve to the data³, the parameters of which describe the level of stigmatisation ultimately obtained and how rapidly it was obtained. The curve used was:

$$\text{stigmatisation} = (\text{lambda} + 0.9) * ((1 - \exp(-\text{time} * \text{beta})) - 0.9$$

This equation is commonly-used to describe natural processes. Lambda describes the point the ‘stigmatisation’ converges towards, while beta indicates how quickly the change occurs. The value 0.9 that occurs in the equation takes into account that the destigmatisation process begins with -0.9. In the documentation of the simulation results the value 'lambda' is

³ Fitting used MATLABs "fminsearch" function.

referred to as "final stigmatisation"⁴. The value given by 'beta x lambda' is referred to as the "speed of destigmatisation"; Beta alone characterises how quickly lambda is reached, so both are used to establish an absolute rate of change.




Figure 2.7: Stigmatisation over time for different levels of common goals.

⁴ Technically 'final stigmatisation' is an inaccurate label, as lambda is never reached, it is the value that stigmatisation tends towards as time tends towards infinity. However as by the end of each simulation stigmatisation is very close to lambda this definition serves for all practical purposes.

Results

The results section opens with a description of a curve fitting method used to interpret the pattern of stigmatisation over time. This is followed by the results from each of the simulations with the individual moderators and finally the interaction between the moderators. There is no need to discuss the effects of neighbourhood as a parameter as the results are the same as in Study Three, however interactions involving neighbourhood are examined to determine whether each moderator influenced the contact effect.

The influence of each moderator on either speed or final stigma was analysed with a two-way ANVOA (Analysis of Variance) with neighbourhood size as second factor. The interaction analysis was conducted with a regression analysis using z-transformed predictors (speed and final stigma). This type of analysis was chosen as it allowed us to gauge the importance of the (potentially significant) 31 interactions between the moderators thus directing the discussion to important simulation findings.

General Destigmatisation Behaviour

There is an observed pattern of destigmatisation in all of the simulations, except those in which the "status" or "contact" parameters were initialized as zero. In simulations with the zero parameter setting, the majority group agents' final opinion was $M = -0.898$ ($SD = 0.0035$), which is not significantly different to its initial setting of -0.90 . These special cases warrant further discussion in their respective parameter's sections, but in general these results are excluded from the analysis as special cases.

Status

There was a main effect of *status* on level of final stigmatisation, $F(2,108) = 45,106$, $p < 0.01$. An interaction between the level of intergroup contact and final stigmatisation also occurs, $F(6,108) = 367$, $p < 0.01$, indicating that status is moderating the contact effect (i.e. the relationship between amount of contact and stigma changed as status changed) These results are displayed in figure 2.8a below.

There was a significant main effect of status on the speed of destigmatisation, $F(2,108) = 14,953$, $p < 0.01$. Here, too, there was an interaction effect with status x intergroup contact on speed of destigmatisation, $F(6,108) = 2,160$, $p < 0.01$. The pattern of results indicated that the more equal status that exists between the agents, the quicker destigmatisation occurs. Additionally the magnitude of the effect increased as contact increases, demonstrating that status operated as a moderator of the contact effect.

The model produced these results because a less equal status led to the majorities (initially negative) opinion of the minority being more important than the minorities (initially positive) opinion in driving intergroup interactions. This causes the interactions to be more negative overall so the majorities' opinion became stable at a much lower level than it would under a condition of equal status. This process also explains the difference in speed, as higher (more equal) status lead to a greater change in average opinion per cycle (as a greater proportion of interactions were positive enough to lead to a positive opinion change), explaining the increased speed of destigmatisation.




Figure 2.8a: Effects of Status on Overall Levels of Stigmatisation.⁵




Figure 2.8b: Effects of Status on Destigmatisation Speed.

⁵ All error bars presented in this thesis show standard errors. As the simulations are taken to be a sample of a larger population they are not corrected for finite population size. Graphs indicating the behaviour of just the agents in the simulation, as if this were the entire population, would have slightly smaller error bars. This would not impact upon the interpretation of the results.

Authority

The second variable tested for direct and moderating effects on level of stigmatisation and the speed of destigmatisation was the influence of authority (implemented as authority agents). There were both main effects of authority on the level of stigmatisation of the minority group, $F(3,144) = 890$, $p < 0.01$, as well as an interaction effect between authority x contact on the final level of stigmatisation, $F(9,144) = 41.5$, $p < 0.01$. Increases in sanctioning of contact by authority agents led to lower overall stigmatisation for the minority group (See figure 2.8b). The effect was larger with more intergroup contact. Authority produced a similar improvement in speed ($F(3,144) = 2,352$, $p < 0.01$) which again can be seen to improve with a greater amount of contact as the authority x contact interaction shows ($F(9,144) = 352$, $p < 0.01$, also see figure 2.10).

The model produced this result because the authority sanctions leads to more interactions per cycle, which in turn allowed destigmatisation to take place more quickly - however this did not affect how much stigmatisation occurred. This effect saturates because an interaction sanctioned by two authority agents is no different to an interaction sanctioned by a single authority agent, thus as more authority agents are added each additional agent becomes less effective, as its influence is more likely to overlap an existing authority agents influence. The change in final stigmatisation occurred because replacing existing majority agents with authority agents leads to a lower proportion of majority agents being in contact with minority agents, producing an increase in final stigmatisation in a similar way to a group size decrease in Study One. Hence, the effect of authorities on final stigmatisation is spurious.




Figure 2.9a: Effects of Authority on Overall Levels of Stigmatisation.

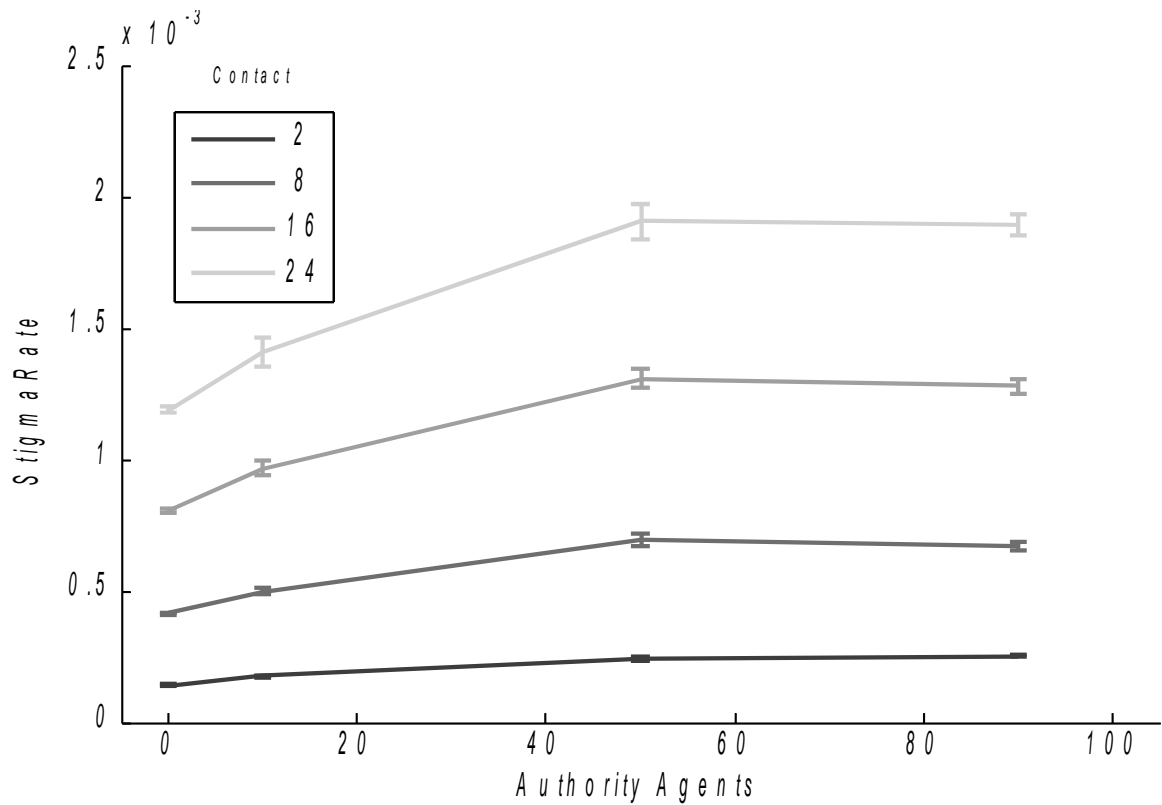


Figure 2.9b: Effects of Authority of Destigmatisation speed.

Cooperation

Cooperation had no significant effect on stigmatisation ($F(3,144) = 2.45, p = 0.066$), nor was there a significant interaction between cooperation and contact on this variable ($F(9,144) = 0.41, p = 0.930$). However, cooperation did have a significant effect on speed ($F(3,144) = 6772, p < 0.01$) and an interaction with contact ($F(9,144) = 988, p < 0.01$). Speed increased with cooperation and the effect was larger at higher levels of contact (see figure 2.12). As existing empirical research has focused on stigmatisation at a particular time point rather than looking at the amount of change over time, it is possible that this increased speed has been mistaken for an increase in final destigmatisation (see the general discussion for

further discussions).

The lack of an effect on final stigma is due to cooperation influencing the magnitude of changes in opinion, but without influencing the direction of those changes. This leads to stigmatisation stabilising at the same point it would in the absence of cooperation, but achieving it more quickly through larger changes in opinion each cycle.




Figure 2.10a: Effects of Cooperation on Overall Levels of Stigmatisation.




Figure 2.10b: Effects of Cooperation on Destigmatisation Speed.

Common Goals

Common Goals had a very similar effect to cooperation. Once again, there was no main effect ($F(3,144) = 1.85, p = 0.141$) or interaction with contact ($F(9,144) = 1.65, p = 0.106$) on stigmatisation. However, a main effect ($F(3,144) = 2269, p < 0.01$) and interaction with contact ($F(9,144) = 341, p < 0.01$) was found on the speed of destigmatisation. An examination of this relationship told a very similar story to cooperation. In this case, the common goals variable leads to an increase in the frequency of intergroup interactions, which in turn increases the speed at which destigmatisation occurred. However increasing the frequency of interactions had no effect upon the magnitude of destigmatisation that was

obtained as the minority and majority opinions still converged to the same point.




Figure 2.11a: Effects of Common Interests on Overall Levels of Stigmatisation.




Figure 2.11b: Effects of Common Interests on Destigmatisation Speed.

Interaction Analysis

These results show that DSIM able to replicate the contact effect, including Allports (1945) "optimal conditions". The significant coefficients for the regression analysis on final stigma are expressed in Table 2.3.

Table 2.3: Regression Analysis of Final Stigmatisation (only significant coefficients are listed).

Variable	Coefficient for stigmatisation
Authority	.162**

Contact	.350**
Status	.888**
Authority * Contact	.084**
Authority * Cooperation	-.007**
Authority * Status	.060**
Contact * Cooperation	.006*
Contact * Status	.220**
Authority * Contact * Status	.039**
Authority * Cooperation * Status	-.006*

* $p < 0.05$ ** $p < 0.01$

The terms with the largest coefficients were status, contact and the interaction between status and contact. This is consistent with the individual findings on the effect of status on stigmatisation described above. The next largest terms were associated with authority and the interaction between contact and authority. Again these are positive terms that validate the findings from the individual analysis, however it is important to remember that the influence of authority is a result of the placement of authority agents on the grid rather than their sanctioning interactions. As such the effect of authority on stigma does not necessarily simulate the effects of authority as Allport (1954) defined it.

The results of the regression analysis on speed of destigmatisation showed that all of the optimal conditions had an effect, as expressed in Table 2.4.

Table 2.4: Regression Analysis for Destigmatisation Speed

Variable	Coefficient for Speed
Authority	.228**
Contact	.730**
Cooperation	.251**
Common Goals	.105**
Status	.349**
Authority * Contact	.192**
Authority * Cooperation	.058**
Authority * Common Goals	-.005
Authority * Status	.091**
Contact * Cooperation	.212**
Contact * Common Goals	.088**
Contact * Status	.293**
Cooperation * Common Goals	.028**
Cooperation * Status	.100**
Common Goals * Status	.041**
Authority * Contact * Cooperation	.048**
Authority * Contact * Common Goals	-.007*
Authority * Contact * Status	.075**
Authority * Cooperation * Common Goals	-.002
Authority * Cooperation * Status	.021**
Authority * Common Goals * Status	-.004*
Contact * Cooperation * Common Goals	.023**
Contact * Cooperation * Status	.083**
Contact * Common Goals * Status	.034**
Cooperation * Common Goals * Status	.010**
Authority * Contact * Cooperation * Common Goals	-.003
Authority * Contact * Cooperation * Status	.017**
Authority * Contact * Common Goals * Status	-.005
Authority * Cooperation * Common Goals * Status	-.002
Contact * Cooperation * Common Goals * Status	.008**

** p < 0.01 * p < 0.05

The regression analysis for speed shows that every parameter and most combinations

of parameters in the model has an effect on how quickly destigmatisation occurs. The largest coefficients were for contact, authority, status and the interactions between these terms.

However all of the coefficients for the moderators and their interactions with contact were large enough to have a positive effect. It is notable that these coefficients were positive, as this is consistent with the results of the individual analysis above, showing that these findings are stable in the face of changes to other parameters.

The vast majority of combinations of parameters produced a positive effect, showing that under most circumstances the optimal conditions facilitate each other's positive effects, which would explain the larger effects observed when the optimal conditions are taken together rather than as individual factors. The highest coefficient for a negative interaction was between contact, common groups and authorities, indeed most of the negative interactions involve these two optimal conditions. As both common groups and authorities are implemented as increasing the number of interactions this is not surprising, as an increase in the number of intergroup interactions has a non-linear effect on stigmatisation whereby each additional increase in the number of interactions produces a smaller improvement (see Study One).

In DSIM units of time are arbitrary, but as a simulation typically involves millions of interactions each can be said to represent a long time. By contrast existing research often only examines how stigmatisation changes over a short period of time. As such it is not possible to compare stigmatisation or speed scores from DSIM to other research, as an increase in speed could be mistaken for an increase in destigmatisation. However DSIM does highlight some parameters as being more important than others and this can be viewed in terms of the existing literature.

Allport (1954), emphasised status, Pettigrew and Tropp (2006) predicted a greater role for authorities (though they conceptualize the conditions as functioning together rather than as individual factors) and these are the two most important factors in the model. Of the four optimal conditions status has the largest effect on both destigmatisation and speed. If we discount authorities' effect upon stigmatisation as an artefact of its implementation, for reasons discussed earlier, it still has a larger effect than the remaining two optimal conditions on speed.

In DSIM Allports (1954) "optimal conditions" all facilitate the contact effect to some degree, with greater improvements in destigmatisation exhibited by conditions previously identified as important, status and authority.

General Discussion

This paper aimed to address a number of important gaps in the intergroup literature, which if answered, shed light on long-standing points of scientific debate. We sought to examine mechanisms involved in the reduction of bias and stigmatisation that might guide interactions (between minority and majority groups). This mechanism implemented in this agent-based model draws on the self-fulfilling prophecy (SFP)-theory of social interactions. In Study One we showed that the model mimics Allport's seminal work on intergroup contact. In addition, the simulations predicated that stigma may be negatively correlated with minority group size due to the fact that while the minority group increases the majority group agents have more and more contact with minority group agents. Study Two empirically confirmed

this prediction in a survey, validating the model. Study Three extended DSIM by using the distance at which agents could interact with each other to manipulate contact. In line with Allport's (1954) contact theory the stigmatisation of DSIM's minority group agents decreased with increasing amount of contact. However and interestingly this effect saturates. Study 4 extended DSIM further by including implementations of the Allport's optimal conditions, equal status between groups, common goals, intergroup cooperation, and the support of authorities. On the whole DSIM successfully replicates the enhanced destigmatisation effect of these conditions postulated in Allport (1954). However the results indicate a more complex relationship between destigmatisation and the optimal conditions as well as interactions between these conditions.

It is surprising that the model shows self fulfilling prophecy (SFP) based behaviour leading to a process of destigmatisation. Typically SFP is assumed to maintain stigma as the circulate processes in SPF are assumed to reaffirm intergroup prejudices (Jussim et al. 2000). The simulations with DSIM indicate that it is crucial for destigmatisation to occur that stigmatised minority groups hold a higher opinion of the majority group. In contrast, if the two groups have a similar opinion of each other opinions don't change and stigma is maintained. It is important to note that the asymmetry between opinions is often found in stigmatised groups as sometimes a group will react to being stigmatised in a manner other than reducing their opinions of the stigmatising group (e.g. Lakhani 2008, Slay and Smith 2011 and Frost 2011) Hence DSIM's starting point is justifiable. Of course DSIM is in many ways a simplification of real processes and these shortcomings of DSIM are addressed at the end of this discussion.

Furthermore it is important to note that destigmatisation only takes place when agents

interact with each other, i.e. there is contact between social groups. The results also replicate experimental findings that show destigmatisation increases with the amount of contact. DSIM operationalises the amount of contact by increasing the neighbourhood in which agents interact (Study 3). Interestingly this implementation results in a saturation effect based on the fact that clusters of agents with similar opinions emerge and at high levels of contact a very small number of very large clusters dominate the simulation. This saturation effect is based on the interaction between neighbourhood size and the density of minority agents, as the neighbourhood size must be sufficient to bridge 'gaps' between minority agents to form clusters of uniform opinion. In DSIM the random placement of agents lead to the minority having a relatively uniform density, but the model predicts that in a situation in which the minority agents were clustered in parts of the environment different results would emerge. In this instance an increase in group size would not lead to significant extra destigmatisation, as they would only influence nearby majority agents, leaving most of the majority with unchanged opinions. Hence we would anticipate that socially or geographically segregated populations may not become less stigmatised with increasing population size.

The final study equipped DSIM with intuitively plausible operationalisations of the four Allport moderators. The status of the minority group was implemented as an influence on the relative importance of the minorities opinion in the interaction, e.g. the higher the status the more the minority group can affect the quality of the interaction. In contrast, common goals and authorities increased the frequency with which individuals were willing to interact, individuals who shared common goals or who were compelled to interaction by authorities that they respected interacted more frequency. Finally cooperation intensified interactions, raising the amount of change, positive or negative, that occurred as a result of interactions.

The simulations with DSIM show that all moderators strengthen the destigmatisation process set-up by the SFP process. Moreover the moderators differentially influence two characteristics of this process, the speed of the destigmatisation process and the final level of stigmatisation. Hereby it is important to note that typical empirical studies would identify both properties as resulting in less stigmatisation. In order to tease these two characteristics apart it would be necessary to examine this in the context of a longitudinal study.

Status affected both properties of the destigmatisation process and also facilitated the effect of contact on the final level of stigmatisation. In contrast the other moderators primarily affect the speed of destigmatisation. The importance of the status results from the higher influence of the minority group on the quality of the interaction at same time leads to more positive interactions as the minority group has a positive opinion of the majority group. Authority had a similar effect in DSIM, but this was not due to the desired implementation (Authorities sanctioning interactions) but was instead caused by a secondary effect of that implementation (Authorities occupy spaces on the grid, changing the ratio of minority and majority agents). This is interesting as it points towards a moderating role for any factor that might influence the ratio of minority and majority members in a given social network, but does not indicate a key role for authority sanctioned interactions.

The finding that other moderators only affect the speed of destigmatisation results from implementations that increase the quantity of interactions rather than changing the nature of them. Existing research (e.g. Schwartz & Simmons, 2001) has noted that the quality of intergroup interactions is a much greater predictor of positive change than the quantity of intergroup interactions.

Another interesting finding is that while Allport's optimal conditions generally

reinforce each others positive effects on speed, there are some conditions under which this does not occur (see table 2.4). As previous research has not examined speed of destigmatisation this prediction cannot currently be verified. However it is important to examine this point further, as most recent research into Allport's (1954) optimal conditions look only at the presence or absence of these conditions collectively.

The literature on intergroup contact, while well-established, fails to test Allport's four moderators of intergroup contact (i.e., equal status between groups, common goals, intergroup cooperation, and the support of authorities) individually and in conjunction within the context of a single longitudinal study. Therefore, it is difficult to determine the relative contributions of these factors, their interactions with one another over time, and contact outcomes. This chapter redresses this gap. The model's general outcome is consistent with the experimental evidence (i.e., intergroup contact facilitates a decrease in intergroup bias and stigmatisation); and explores the unique contribution of each moderator to overall reductions in bias over in the longer-term.

Limitation & Outlook

This chapter presented for the first time is an agent-based model of Allport's (1954) contact hypothesis of intergroup contact. The majority of empirical work has focused on establishing the validity of this hypothesis and positive impact of the Allport's four moderator. There has also been work revealing numerous other moderator facilitating destigmatisation, e.g. impact of personality factors (Hodson 2008); intergroup anxiety (Turner et al. 2008) and

imagined intergroup contact (Crisp and Turner 2009) to mention a few. Against this background it is clear that DSIM captures only a very limited number of factors. Nevertheless the chapter illustrates that an agent-based approach can produce empirically valid results and lead to novel theoretically relevant findings. Furthermore the agent-based approaches present an ideal framework for future research into intergroup contact as this framework allows modellers to integrate crucial processes in an intuitive way. For instance relevant cognitive processes into agents, e.g. imaginary contact, can be integrated into the agents' behaviour. Social processes can be added to the agents' interaction, e.g. communications about extended contacts. Finally (but without being exhaustive) agents' mobility may mimic effects of geographical or social space on destigmatisation. For instance, Binder et al. (2009) showed that less prejudiced individuals to be more likely seek contact with stigmatised groups. In this context it is also worth mentioning that Binder et al. (2009) touch on an important issue in the contact literature, the causal direction between contact and reduced stigmatisation. Here agent-based models can also contribute to the debate as the direction of causation in an ABM is always clear from the implementation of the model.

This chapter as a whole underlined the fruitfulness of the agent based modelling approach in tackling problems in social psychology. The approach is strongest when it is applied to situations involving group properties emerging from a great number of smaller social interactions, the thesis takes this line of research forward using a more general model of group dynamics before going on to look at ways to improve how the agent based modelling method is used in social psychology.

Chapter Three: Modelling Fixed Social Group Sizes

The model presented in Chapter 2 investigated contact theory by providing a detailed view of the implications of previously identified moderating factors. While the model was a success, the agents were static in most important ways, they did not change their sphere of social influence or group membership. This chapter presents a more dynamic model of social group behaviour, focusing on social group size, after the survey study in the previous chapter highlighted the importance of this variable.

Social Network Size

Human social networks are an important area of study for social scientists in several disciplines (see Wasserman & Galaskiewicz, 1994, for a review). One facet of this is the number of connections that a given individual is capable of forming. Researchers have been examining the size of human social networks for over two decades, but findings in this area have been inconsistent. Researchers have used various methods to estimate network size, McCarty et al. (2001) finds a network size of 291, using two different methods, both of which estimate total network size by asking a participant how many individuals they know in a given sub-population and calculating the size of their overall network using the national frequency of individuals within that sub-population. By contrast Pool and Kochen (1978) examined several methods, such as developing a model based upon the number of new unique acquaintances an individual contacts each day which produced estimates ranging from 1,100 to 4,250 individuals.

Hill and Dunbar (2002) stand out in this field by providing two important contributions. The first is that they identify that the discrepancy in size estimates are a result of inconsistencies between how researchers measure social networks. They go on to look at the size of surviving tribal societies and the size of individuals Christmas card lists. They argue that this is a better estimate of their participants current valued number of social connections than existing methods as it shows that participants are willing to make at least some effort to maintain their social connections with these individuals as opposed to merely being able to recall their names. The second is that they explain why human social networks should have a size limit in the first place. Dunbar (1993) studied primate social network structure and brain structure, finding that neocortex size is a good predictor of social network size. Applying these observations to human neocortex size, he predicted a human social network size of 147.8 individuals. Hill and Dunbars (2002) test supports this prediction, finding that the average number of Christmas cards an individual sends, 153.5. They conclude that neocortex size is a good predictor of the size of human social networks, which provides evidence for social network sizes having a biological basis. This line of research has been extended to look at other factors such as perspective taking ability (Stiller & Dunbar, 2007) and personality factors (Roberts, Wilson, Fedurek & Dunbar, 2007) and is summarised in Dunbar (2008) which continues to emphasise the biological account.

The line of research described in the previous paragraph examines the size of an individual's social network, but has increasingly been used to discuss the size of social groups. The distinction is that an individual's social network consists of people who are directly connected to the individual, a social group consists of everyone who identifies with that group, (see social identification theory in the introduction) regardless of the personal relationships. The biological account could not explain a limit of social group size, as an

individual's limited capacity to form connections would not stop unrelated individuals from starting to identify with the group. Dunbar (2008) highlights situations in which social groups with roughly 150 members have occurred, arguing that this is due to human social network sizes restricting some groups to the level at which everyone within the group can know each other personally. In the wider media the concept is applied more broadly to social groups, with social networks and social groups often used interchangeably (e.g. Allen, 2004)⁶. The goal of this chapter is to investigate factors that influence the size of social groups, in order to see if the concept of a fixed group size is applicable and to offer new insights into the reasons that fixed social network sizes are observed, this will be achieved using an agent based model to create a test a model of group dynamics.

Social Psychology and Group Size

Developing a theory of social group size requires a model of the processes that influence a social group's size. Such processes will describe the manner in which individuals interact with social groups. Considerable research interest has been focused on how individuals behave in relation to social groups, much of which has been driven by social identity theory (Tajfel & Turner, 1979; Turner, 1987; Hogg and Abrams, 1998; Abrams and Hogg, 2000; Jetten & Postmes, 2006)

Social identity theory. Individuals maintain a number of social identities. These are social groups that a person considers themselves to belong to or that become salient ways of categorising individuals. The characteristics of these groups can vary dramatically as an

⁶ This blog is the second hit when googling "Dunbars Number" at the time of writing (The first being the wikipedia entry). Many of the subsequent entries observe the same pattern, Dunbars number is applied to group sizes rather than social network sizes.

individual can form an identification based on a number of characteristics. Such examples might include: identities with groups as diverse as “English” “Christians” or “Members of Hill Road Social Club”. Any individual that considers themselves to belong to a social group will develop a social identity based upon that group. However, social identity theory (Tajfel & Turner, 1979) holds that individuals do not value their social identities equally. For instance an individual might identify more strongly as a student than they do as a gambler and will treat these identities differently. Specifically, high identifiers will go to lengths to emphasise that group membership and conform to that groups norms, whereas one who has a very weak identification with the target group will not consider it as important to defining sense of self, and may seek to leave that group (Sani & Pugliese, 2008).

Also, it is important to highlight that identification is not static but dynamic. Therefore, an individual's identification with a social group can change over time. Sense of identification can be strengthened through positive experiences, experiences in which the target identity is particularly salient. On the other hand, negative experiences can lead to a reduction in identity, notably social rejection from another group member (The strong impact of ingroup rejection is highlighted by Crocker & Major, 1989, who note that common methods of preventing esteem loss following negative interactions are not effect in interactions with the ingroup). Also, the mere anticipation of negative experiences from within the ingroup can also lead to reduction in identification and loyalty (Jetten, Branscombe, Spears & McKimme, 2003). However, not all social rejection will adversely affect social identification, the rejection-identification hypothesis (Branscombe, Schmitt & Harvey, 1999) holds that rejection from an individual who is not a member of the social group will strengthen, rather than weaken, the identity in question. This holds true for many different types of groups ranging from large groups that individuals are born into, such as

gender or ethnicity (e.g. McKoy & Major 2003) to groups that people join voluntarily, such as people with piercings (Jetten, Branscombe, Schmitt & Spears, 2001).

Losing identities. As mentioned above, identification changes are important to a model of group size as they influence when individuals will leave a social group. Most researchers find that low identification causes individuals to leave groups as a matter of course, for instance Sani and Pugliese (2008) find that low identification is always associated with a high intention for forming a schism, but that other factors moderate how strong this effect is. In order to produce a model of the dynamics that determine a social group's size, this chapter will focus on an individual's level of identification with a group as a determining factor in whether that individual will choose to belong to that group. In the model when individuals have positive experiences in the context of their group membership then their identification with that group will increase, but if they have a negative experience with an in-group member then it will decrease. While in-group rejection does not universally lead to lower group participation (e.g. Williams & Sommer, 1997, find that rejection increases social loafing in men but not in women) it leads to identification loss frequently enough that ingroup rejection can be modelled as causing a reduction in identification (e.g. Cowan & Ullman, 2006, study the impact ingroup rejection in female collage students).

Leaving a social identity. As mentioned above, identification changes are important to a model of group size, as they influence when individuals will leave a social group. Most researchers find that low identification with a social group causes individuals to leave groups. For instance, Sani and Pugliese (2008) find that low identification is always associated with a high intention for forming a schism, but that other factors moderate how strong this effect is. In order to produce a model of the dynamics that determine a social group's size, this chapter will focus on an individual's level of identification with a group as a determining factor in

retaining group membership. In the model, when individuals have positive experiences - in the context of their group membership - then their identification with that group increases. However, if the agents have negative experiences with in-group members, then identification with the group decreases. While in-group rejection does not universally lead to lower group participation (e.g. Williams & Sommer, 1997, find that under some conditions rejection increases social loafing rather reducing identification) it leads to identification loss frequently enough that ingroup rejection can be modelled as causing a reduction in identification (e.g. Cowan & Ullman, 2006).

Forming new identities. A model of social group size will need to model the processes that cause individuals to join social groups, as well as those that motivate individuals to leave them. The research into identification changes is only relevant where an individual already identifies with the group. The process by which an individual joins a new group will be modelled using conformity research. It is widely accepted that individuals can join a social group through conformity; the degree of social power exhibited by a group increases as it gains more members (Latane, 1981). Conformity can lead to social conversion, for instance (Nail & Helton, 1999) investigate compliance and conversion as forms of conformity, while much of the work in this area focuses on conformity of beliefs and attitudes, having a social identity is strongly related to sharing the beliefs and attitudes of the group (Tajfel & Turner, 1979). The role of conformity has been widely examined, Asch's (1955) seminal conformity experiments continue to be investigated (see Bond, 2005, for a meta-analysis relating to conformity group sizes) so there is a strong research base to draw on for implementing conformity. Additionally, existing models of group dynamics have successfully used conformity to model other processes, such as in Iannaccone & Makowsky's (2007) model of religious regionalism.

The Current Research

Existing research on social network size has concerned itself primarily with biological factors (e.g. Hill & Dunbar, 2002, Dunbar, 2008) or cognitive abilities (e.g. Roberts, Wilson, Fedurek & Dunbar, 2007). The goal of the present research is to investigate whether a group of a fixed size could be formed through social factors, notably social conformity and in-group rejection. If such a process is observed it provides a psychological account of fixed group sizes and will also suggest a psychological explanation of social network sizes, complementing the existing biological explanations.

This will be achieved by developing an agent based model to simulate the effects of social conformity and ingroup rejection upon a theoretical group. In addition to examining whether groups of fixed size emerge, the model will also show the number and distribution of agents that emerge under different parameters. This enables it to address more general questions about the behaviour of social groups, such as investigating the conditions that cause a social group to propagate through a society, acquiring individuals who were far removed from its existing members.

The rest of this chapter is organised as follows, firstly it describes the dynamic groups model alongside the parameter settings used in the investigation and the outcomes measured. Then the results section describes the most common patterns observed in specific simulations before tackling the general case results for each of the outcomes measured. Finally, the chapter concludes by discussing the consequences of these outcomes in the context of the literature discussed above.

Method

Dynamic Groups Model (DGM)

The DGM consists of a 51 by 51 grid, each space of which is either empty or occupied by a single agent. The grid is an abstraction of the relationships between the agents that occupy it. Agents that are close together on the grid have strong social ties and interact frequently, whereas agents that are far apart on the grid do not interact with each other frequently. Agents can alter the agents that they interact with by changing their position on the grid. Each simulation is initialised with agents occupying 70% of the grid spaces, thus simulations are initialised with 1,821 agents⁷.

Only one social group is explicitly modelled, which each agent may or may not belong to. As no rule in the simulation allows an additional social group to be formed, there will never be more than one social group. On a theoretical level these agents represent individuals who belong to several social groups, but the additional groups are not represented within the simulation. Consistent with social identification theory, each agent has a level of identification with the simulated social group. This is represented by a number ranging from 0 (Does not identify with the group) to 1 (Strongly identifies with the group). Agents possess no properties other than their position on the grid and their level of identification with the group. Agents with an identification of 0 do not belong to the group and are termed 'nongroup agents'; agents

⁷ There is nothing special about this number, while grid size was not varied extensively in this model pilot studies did not reveal a high sensitivity to grid size.

with an identification above 0 are members of the group and are termed 'group agents'.

At the start of each simulation 140 of the agents in a 20 by 20 square on the grid immediately increase their identification with the group to 0.001. This means that approximately half of the agents in this area will weakly identify with the group at the start of the simulation, the other agents in the simulation will all be nongroup agents. The starting location for group agents is limited to a 20 by 20 square in order to study the properties that a social group might develop. Initially the group is coherent, (Most group agents are interacting with at least some other group agents) integrated (i.e. Most group agents are interacting with at least some non-group agents) and not propagating (There are no group agents outside of the area in which the group is initialised). The model would not be able to study all of these properties if group agents were initially common over the whole grid. In a pilot study agents were initialised as belonging to the group throughout the whole grid. The behaviours of agents under these conditions were not qualitatively different to those who do not leave the 20 by 20 square in the final simulation.

When two agents belonging to the group interact it must be determined whether ingroup rejection occurs. This is implemented by classifying each interaction as either a “positive interaction” in which no rejection occurs, or a “negative interaction” in which it does. Under different circumstances social groups can experience different levels of in-group rejection (e.g. Cowan & Ullman, 2006) so the level of positive interactions will be varied between simulations. If rejection occurs (a negative interaction) the agents involved reduce their identification with the group, otherwise a positive interaction has taken place and the agents identification with the group is increased. Consistent with the literature (e.g. Sani & Pugliese, 2008), a decrease in identification may cause an agent to leave the group. The parameter describing the degree of in-group rejection shall be termed "positive interactions"

and it shall describe what proportion of interactions is positive.

If random number > parameter then both agents decrease identification

If random number <= parameter then both agents increase identification

When a group agent interacts with a nongroup agent the chance of a positive interaction is not calculated. The rejection-identification hypothesis (Branscombe, Schmitt & Harvey, 1999) states that out-group rejection results in increased identification with the ingroup. As such when a group agent interacts with a nongroup agent the group agent will always increase its identification, as either it would have a positive interaction and increase its identification with a valued social identity or it would have a negative interaction and increase its identification as would be predicted by the rejection-identification hypothesis.

When two nongroup agents interact, the interaction has no effect, as in the absence of a group agent the interaction will have no bearing on identification with the target group. Also, the nongroup agent does not change its identity in the group-nongroup interaction described above. This is because they do not identify with the group and thus cannot gain identification with the group through having it as a salient valued identity or through having it rejected by the outgroup. Instead nongroup agents can start to identify with the group through conformity.

Modelling social conformity as a force that causes individuals to adopt a group membership on a temporary or permanent basis is a tried and tested approach from other psychological computer models (Iannaccone & Makowsky, 2007). In DGM it is implemented in the following manner: Each cycle any nongroup agent that has more group neighbours than

nongroup neighbours has a one percent of joining the group.

*If number of group neighbours > number of nongroup neighbours and random(100) = 1
then join group.*

While conformity has a low chance of occurring on any particular cycle, the test is repeated every cycle that the agent continues to be outnumbered. This means that the overall probability of an agent joining the group is dependent upon how long the situation in which they are outnumbered by group agents persists. The conditions for conformity persist longer if the number of group agents adjacent to the nongroup agent is larger, as more of the group agents could move away without ending the condition. This implicitly implements the finding that conformity is more likely where the majority is larger (e.g. Asch, 1955) and that larger groups have more social power (e.g. Latane 1981).

Joining a group is implemented as a positive identification change, the agents identification is increased as if it were a group agent having a positive interaction, which results in it obtaining a level of social identification above zero. This causes the agent to become a group agent. Identification changes are handled using the learning algorithm introduced in the previous chapter:

*To increase identity: $New ID = (Old ID + 0.001) - (OldID * 0.001)$*

*To decrease identity: $New ID = (Old ID - 0.001) - (OldID * 0.001)$*

If the resulting identification is below zero, then the agents identity is set to zero and it

becomes a nongroup agent. Once individuals do not consider themselves part of a group they cannot identify less with that group, so negative identification is not possible. This equation is used because it avoids polarising the agents' final identification scores. If interactions changed identification in a linear manner then an agent who increased its identity during 51% of the interactions would have the same identification as an agent who increased its identity during 100% of interactions. It is helpful to be able to distinguish between these cases, so a more nuanced update rule is adopted.

In order to determine when an agent will decide to change its position on the grid in order to interact with different individuals, a preference for similar individuals is implemented, as an individual's preference for similar individuals has been well established (e.g. Fowler, Settle and Christakis, 2010). Existing implementations of this concept have proven successful in spatial agent based modelling (e.g. Pralat & Zhang, 2011), so there is solid evidence for thinking that it would work well with social psychological constructions of space. In DGM this is implemented as follows: Agents have the opportunity to move at the end of each cycle after interacting with their neighbours. Each agent makes a decision about whether it is satisfied with its current position. It does this by comparing the ratio of agents similar to it (in terms of whether they belong to the group or not) and the number of agents that are different to it, moving if the ratio is too low. The level of different agents that an individual is willing to accept is the second parameter manipulated in DGM experiments. It is termed "similarity preference" as it describes the degree to which agents prefer to be with similar agents.

If similar agents / (similar agents + different agents) < similarity preference then move

Simulations

Simulations were performed to look at the ways in which agent level interactions, driven by the factors implemented above (conformity, ingroup rejection, rejection-identification), might influence a social groups size and structure. The object of the simulations is to determine when and how a group with a fluid membership but fixed size might emerge. In each simulation a small group of agents were initialised as belonging to the group (initially low identification) and the majority were initialised as not belonging to the group (zero identification). The rationale for using a low identification for the agents that belonged to the group was to avoid making the number of agents initially belonging to the group too influential, as a low identification allowed the agents to leave the group easily if the dynamics of the simulation supported this behaviour.

The parameter "positive interactions", which determined the percentage of in-group interactions that would be positive, was varied between 0 and 100 in increments of 10. The parameter "similarity preference", which determined what percentage of adjacent agents needed to be similar to an agent in order for it not to move, was varied within the same range. A simulation for each combination of parameters was performed 10 times. These two parameters were the only independent variables systematically manipulated throughout the DGM simulations.

Data Analysis

To examine the emerging group dynamic five dependant variables were observed. The social group's size, the average ratio of group to nongroup neighbours each group agent had, the average identification of group agents with the group, the number of group agents outside of the groups initial area and the number of agents that had changed their position in the social space during the final ten cycles of the simulation.

The final group size measure was taken in order to test the primary hypothesis: That the level of in-group rejection in a social group acts as a limit on the size of the group. The ratio of group to nongroup agents was a measure of how coherent the group was, as this value becomes higher when more group agents are in contact with each other. This measure is used to ensure that the simulation was still producing realistic social groups that had members interacting with each other where in-group rejection was strong enough to limit the size of the group.

The average identification of all group agents was taken in order to explore the conditions that might lead to groups with more or fewer strongly identifying members. The fourth outcome, how many group agents existed outside of the groups initial area, enables an assessment of the conditions in which a small group can proliferate through a social network and obtain members that initially had no contact with their group.

The final measure, 'how many agents moved in the last ten cycles', was a measure of the stability of the simulation. In pilot studies, I found that a simulation would not stabilise, with most the agents moving every cycle. As most of the above measures are obtained during the final cycle of the simulations execution, they become unreliable if that cycle is not a

typical state of the simulation. Thus, a final measure to identify simulations that had not reached a stable state was required. If at the end of a simulation more than half of the agents in the simulation had moved in each of the final 10 cycles, then that simulation was declared to have a chaotic outcome and its results were not used in any other analysis.

Results

Before examining the outcome variables, a typical simulation is described to provide context in order to understand the model's other results. Then the analysis for each of the dependant variables is presented in turn, these are group size, identification, segregation (as a measure of group coherency) and network propagation. In all cases chaotic simulations have been excluded from the results.

Chaotic Simulations

Under some conditions a "chaotic simulation" emerged in which agents are consistently unable to find a position on the grid that satisfies their similarity preference causing them to constantly change their position (i.e., who they interacted with). As stated above, any simulation in which more than 2% of the agents moved position during the last 10 cycles was identified as chaotic. All of the simulations with a "similarity preference" at or above 90% resulted in a chaotic simulation, but no other parameter or combination of parameters did, thus 990 simulations remain to be analysed.

Example Simulation

The results reported below are from a single simulation in which both parameters were set to 60. While there is no single prototypical simulation, many of the behaviours observed under these parameters will occur to a greater or lesser extent in other simulations. The images (figure 3.1) illustrate the simulation at 0, 1, 50 and 100 cycles. Each dot represents an agent. Group agents are represented by blue dots, with brightness proportional to its level of identification with the group. Nongroup agents are represented by grey dots.

In the simulation's initial state, the agents belonging to the group all have low identification and remain in close proximity. In the first cycle, any agent who is not satisfied with its initial position moves randomly, leading to the agents being distributed widely. Also, this leads to a large number of interactions with the outgroup and which produce high identification for the group agents. However, these group agents continue to move randomly until they find a position on the grid with sufficient other group agents to satisfy their similarity preference, which in this instance required 60% of their neighbours to belong to the same group. While many agents immediately moved in the first cycle, due to their random initial position being unsatisfactory, by cycle 50 they have all returned to the initial area. This is because many of the group members did not leave the initial area so it is likely to be the first location the agents randomly move to that satisfies their preference.




Figure 3.1: Simulation state at cycles 0, 1, 50 and 100

In the simulation's initial state, the agents belonging to the group all have low identification and remain in proximity. In the first cycle any agent not satisfied with its initial position moves randomly, leading to the agents being distributed widely on the first cycle. This also leads to a large number of out-group interactions, producing high identification through rejection-identification. However these agents continue to move randomly until they find a position on the grid surrounded by agents that they want to interact with. While many

agents immediately moved in the first cycle, due to their random initial position being unsatisfactory, by cycle 50 they have all returned to the initial area as the similarity preference parameter encourages them to seek out agents from their group and many did not leave the initial area.

By cycle 50 the agents belonging to the group have formed social links with each other and while the agents have settled into several distinct clusters, though the agents in them are all considered to be members of the same overarching social group. Over the next 50 cycles (50-100), it becomes apparent that these clusters are stable, as the agents' locations and levels of identification do not undergo much change. Interestingly, however, agents with outgroup neighbours have a significantly higher identification than those surround by ingroup agents (Visually, the blue dots with adjacent grey dots are brighter). I will pick up on this in the discussion.

This example simulation will be used to contextualise the results from other simulations. There are some systematic deviations from this example simulation, for instance under parameter settings involving a lower preference for similar neighbours the higher identification for agents in contact with the out-group becomes more pronounced. Most of the agents on the edge of clusters develop high identification while those in the middle have a much lower identification. Under parameter settings involving fewer positive in-group interactions those agents in the centre of the cluster will lose enough identification to dis-identify with the group and leave, producing a ring of agents surrounding the area where a cluster once was. Variations such as these will be addressed as they become relevant to the results.

Group Size

The main hypothesis framing the research was that the size of a small social group would be limited by the psychological factors implemented in the DGM. Importantly, this concerns the total number of agents that identify with the social group, not just those that are in a particular cluster. Figure 3.2 shows the relationship between the models parameters and the group size outcome




Figure 3.2: Effect of 'similarity preference' and 'positive interactions' on group size.

The results show that an increase in the number of positive ingroup interactions, or equally a decrease in the level of similarity preference, lead to the ingroup having larger number of agents. The increase is not linear, so in the absence of a direct effect the data are analysed by fitting a logistic sigmoid curve to the results for each setting of ingroup bias. The mean centre point for the sigmoid curves occurs at a positive interaction level of 55.17. As this is based on simulations that varied this parameter in ten point increments, it is only possible to conclude that the point at which a positive interactions improvement will make the most difference is somewhere between 50 and 60. In other words the closer the proportion of positive interactions is to the critical level (between 50 and 60 percent) the greater the impact that changing the level of positive interactions within the group will ultimately have upon the group's size.

To understand why there is a critical point between 50% and 60% positive interactions, it is necessary to understand the behaviour of clusters of group agents. As discussed earlier, there is a behaviour, shown in figure 3.1, whereby ingroup agents at the centre of a cluster will lose identification due to a high number of negative ingroup interactions. If positive ingroup interactions are low, the agents in the centre of these clusters can feel rejected enough to leave the group entirely. At even moderately low levels, any two group agents in contact with each other will cause one agent to leave the group. If more than 60% of the ingroup interactions are positive, this no longer occurs, instead large clusters of group agents with high identification are 'viable', resulting in higher group sizes. On the other hand if ingroup interactions are primarily negative group agents will not be able to maintain their identities while in contact with each other causing many of them to leave the group (See figure 3.3).




Figure 3.3: Comparison of simulation progression for high (80) positive interactions (top) and low (20) positive interactions (bottom) at average similarity preference (50) at cycles 1, 50 and 200 (left, middle and right respectively).

There is a tipping point, above which large groups are more sustainable through their positivity, which leads to the pattern observed in above (figure 3.2). This also influences the interaction between positive interactions and similarity preference. Specifically, in simulations with a low similarity preference, agents might maintain their identities with the group by interacting with a lot of nongroup agents. But, if similarity preference is high, then they no longer have this option; they will change their social links if they are primarily interacting with nongroup agents, resulting in them being unable to maintain their identity with the group. The result is a very small group - as ingroup rejection drives agents to leave the group

while a high similarity preference makes ingroup rejection more common.

Finally, high similarity preference is also associated with a smaller group size, even while most ingroup interactions are positive. Again, a high preference for similar agents to interact with leads to fewer intergroup interactions, decreasing the chances of a nongroup agent encountering enough group agents to conform and join the group.

Coherency and Integration

When initialised the ingroup was coherent (i.e., most group agents interacted with in group agents) and integrated (most ingroup agents interacted with outgroup agents). Previous research suggests that individuals interact most with people in ‘similar’ social groups (e.g. Fowler, Settle & Christakis, 2010), which implies that for the findings of a simulation to be valid a level of coherency should be maintained throughout the simulation. The extent to which groups integrate with each other is variable and DGM can help to understand this variance. Coherency and integration can be understood by examining the segregation of the group that results from different parameter settings. The segregation of an individual agent is measured by calculating the ratio of agents adjacent to it; and a segregation score for the group can be obtained in each simulation by averaging these results. (See Figure 3.4)




Figure 3.4: Effects of 'similarity preference' and 'positive interactions' on segregation.

The extent of segregation in the simulations was increased when either similarity preference ($F(8) = 1097.24, p < 0.01$) or positive interactions ($F(10) = 218.15, p < 0.01$) were increased. An interaction is also detected ($F(80) = 57.16, p < 0.01$) which reflects the observation that at higher levels of similarity preference positive interactions had less effect, at the highest level of similarity desire that did not lead to chaos (eighty) the groups was always completely segregated (segregation equals one) regardless of the level of positive interactions.

This was because with sufficiently high levels of similarity preference the simulation stabilised with no contact at all between group and nongroup agents, such that any agents moving into contact would always have too many different neighbours and move again.

Additionally at high levels of ingroup positive interactions adjacent group agents would never lose identification with the group (through ingroup rejection) and any outnumbered nearby non-group agent would eventually be converted, leading to dense clusters of group agents. These factors interact as the isolationism produced by high levels of similarity preference reduces the opportunities for conversions through high positive interactions to take place.

Identification

While I have already highlighted patterns of identification strength for a single simulation, there are two distinct patterns in the relationship between these parameters and the identification outcome worth further note (figure 3.5).




Figure 3.5: Average Identification

At ingroup positive interactions 60 and above there is a very clear pattern, an increase in positive interactions leads to a linear increase in group size. Lower levels of similarity preference also produce slightly higher identification scores. As more positive interactions mean that ingroup interactions are more likely to produce identification increases the first pattern is not surprising. Additionally a lower similarity preference means that group agents are more likely to be in contact with nongroup agents, which leads to extra identification increases through the rejection-identification rule. However a different pattern occurs at lower levels of positive ingroup interactions.

When ingroup positive interactions are less than 60, the results are far less intuitive. An increase in the number of in-group positive interactions leads to a decrease in

identification. To understand this result it is necessary to understand the circumstances in which a group agent with low identification can remain a member of the group. If an agent simply loses identification repeatedly, it will leave the group and have no influence on the average identification of group members. To maintain a low identification without leaving the group, an agent must have two things. Firstly it must have a nearby group agent with which it has predominantly negative interactions, causing it to lose identification. Secondly, it must also have one or more nearby nongroup agents which will cause it to gain identification through rejection-identification. If this second condition is not met the agent will not maintain a low identification, it will simply cease to identify with the group.

The more negative that the interactions within the group are, the greater the amount of outgroup rejection that would be required to cause a group agent to continue to identify with the group. Thus low identification agents become rarer at lower levels of positive in-group interactions, producing a higher average identification with the group. Consistent with this explanation, earlier simulations with a positive interactions parameter of less than 60 often produced very small groups (See figure 3.2).

Network Propagation

The final point of interest is the conditions under which the group propagates through the grid, allowing group agents to persist outside the 20x20 area in which they were initialised. This is measured by counting how many group agents ended the simulation outside of their initial area; these agents are referred to as 'outsiders'. Two dimensional representations of this data do not describe it adequately as the pattern that emerges is that there are two

optimal sets of parameters that produce a noticeable level of outsiders, see figure 3.6.

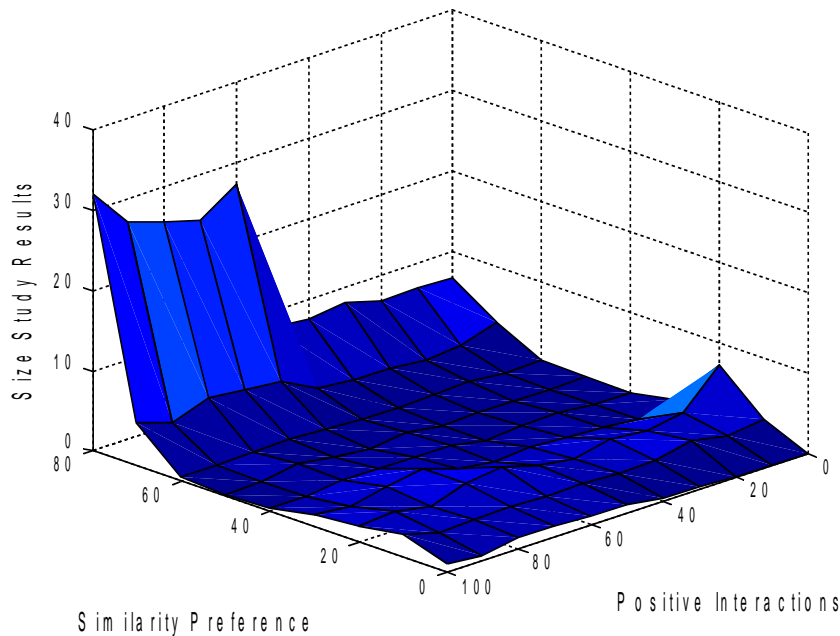


Figure 3.6: Conditions under which the group moves beyond its initial area.

The first occurs at a high level of similarity preference and becomes much more pronounced if a high level of positive interactions are also present. Under these conditions two things happen to create a large group presence outside of their initialised area. Firstly, at the start of the simulation, most group members are not content with their initial neighbours, as on average only fifty percent of them belong to the group. This causes them to move. The majority of agents that move will be unsatisfied with their new position and move again, repeating this pattern of unsatisfied preferences and movement until they return to the starting area. At this stage, the initial group area is mostly populated by group agents as all nongroup agents will have conformed and joined the group, or moved away due to their own unsatisfied

preferences. However, as the agents outside of the initial group area move around, the high level of similarity preference will cause some nongroup agents to move away from the group agents, creating areas containing no agents of either type.

When another group agent moves into the centre of this space they are no longer in contact with any nongroup agents. This allows such group agents to persist outside of their initial area. However if positive interactions are high, other group agents might form clusters with these outsider agents. When the second agent moves to join the group, they create a cluster that is stable owing to the high level of positive interactions (members' identification remains high; and the high local density of ingroup agents makes them resistant to moving). Over time, more agents join the cluster to fill the available space. The agents in these clusters are now unlikely to be dislodged, as outgroup agents that move nearby will likely join the group or move away again, thus preventing enough of them accumulating to influence the small group cluster. (See figure 3.7.).




Figure 3.7: A 'positive interactions 80 similarity preference 80' simulation at 1, 25 and 50 cycles

The second condition leading to higher levels of outsiders in figure 3.6 is illustrated by a smaller peak in the behaviour which demonstrates a more typical behaviour. This occurs under less extreme parameter settings and does not result in the outsiders having no interactions with nongroup agents. Instead this effect occurs at a low level of similarity preference and a low level of positive interactions. In this situation agents will tolerate a high number of outgroup neighbours. Additionally, segregated clusters belonging to the group are unable to survive. This is because the group experiences a higher level of ingroup rejection (low positive interactions) resulting in dense groupings of group agents driving some of their members to leave the group. Agents can only maintain their group identity if they have some interactions with agents outside of the group to reinforce this identity, this makes mixed areas of group and nongroup agents more common as well as increasing the number of agents that leave the initial group area. See figure 3.8 for time steps from a simulation in which this occurs.




Figure 3.8: A 'positive interactions' 20 'similarity preference' 20 simulation at 1, 50 and 200 cycles.

Such simulations never completely stabilise. For example, consider the nongroup agents close to the main cluster at the bottom left of the simulation. For some of them, more than twenty percent of their neighbours are also nongroup agents, so they will not move. They are also outnumbered by group agents so they will eventually conform and join the group. As soon as one joins the group, the relatively high number of group agents adjacent to it will cause it to experience sufficient ingroup rejection to lose identification and leave the group. This pattern will repeat indefinitely. While the simulation will never completely stabilise, it will not change in any substantial way either, so results from such simulations are still valid.

Discussion

The DGM is a model of group dynamics driven by social identity theory (Tajfel & Turner, 1979; Turner, 1987; Hogg and Abrams, 1998; Abrams and Hogg, 2000; Jetten & Postmes, 2006). The identification changes implemented in the model emphasised rejection-identification theory (Branscombe, Schmitt & Harvey, 1999) and social conformity (e.g. Nail & Helton, 1999). The theoretical basis for the DGM centres on the concept of identification, as agents activities are driven by and influence their identifications. The key results of the model are understood in the context of why the changing identification of individual agents produces a particular group-level behaviour, for instance how the size of a group has been influenced by the changing identities of the individuals within it.

The model showed how changes to the group's level of ingroup rejection and agents' preference for similar individuals altered several characteristics of the group. These were the extent to which its members identified with the group, the coherency and integration of the

group (measured by segregation) and the size of the group. The primary goal of the DGM was to explain the variation in group size, and to determine if existing limits of social network size (e.g. Hill & Dunbar, 2002) could be applied to social groups. The results suggest that under most conditions a fixed group size can be obtained. The influence of the positive interaction and similarity preference parameters on group size are discussed in detail below.

Group Size Outcome

Examining the influence of social identification and social conformity on social group size was a primary aim of this model. Dunbar postulates that an individual's social network has a fixed maximum size (e.g. Dunbar, 2008) and that there are individual differences that mean that this limit is not always reached (e.g. Stiller and Dunbar, 2007). Previously whether social network size limits might apply to social groups has not received much formal attention.

In DGM, the extent to which agents experience positive ingroup interactions determines the size of the social group. Additionally, the degree to which group agents seek to interact with their ingroup influences this. That is, groups with a higher tolerance for diversity are more successful at attracting and retaining more agent members. There is a critical level of positive ingroup interactions (around 50-60%) above which groups are much larger, the influence of this critical level is higher at greater degrees of tolerance for different agents, so DGM predicts a tipping point in ingroup rejection which would lead to large changes in group size, especially for tolerant groups. In terms of the main research question, DGM predicts that social groups will have a fixed maximum size primarily driven by their level of ingroup

rejection.

While DGM is a simulation of social group size rather than a simulation of the size of an individual's social networks, there are similarities that allow DGM to inform ongoing investigations into individual's group sizes. A fixed group size is acquired in the DGM because larger groups become more vulnerable to losing members as a larger number of members provides more opportunity for ingroup rejection. An individual's social group may encounter a similar limitation, as an individual has more friends the chance of there being conflicts between their friends that make it possible for them to maintain both relationships increases. Further work would be needed to test this hypothesis. These findings are interpreted as follows - DGM compliments the existing reasons for the finding of a fixed maximum group size (e.g. biological reasons, Hill and Dunbar, 2002) by suggesting that part of the explanation for why some networks do not reach this limit is that the limit is changed by two additional factors. Previous research has focused on the traits of the individuals whose social network is being examined as reasons for the networks size, however DGM looks at a new type of factor by considering the attributes of the members of the network. DGM shows that the properties of the individuals that make up a network influence its size in two ways. Firstly the tolerance for diversity that individuals possess influences the size of networks that they are in, with greater tolerance promoting larger networks. Secondly, DGM suggests that the degree of positive interactions between members of the network will have a strong effect upon the overall size the network can reach, with a greater proportion of positive interactions leading to larger network sizes. The explanation for this is that when networks contain large number of individuals and have a high level of ingroup rejection then these individuals reject each other, causing some to leave and reducing the size of the network. Conversely when there is a high level of tolerance for diversity more network members are in contact with

individuals outside of the network and have more opportunities to recruit additional members.

Identification Outcome

The second main outcome of DGM is the pattern of identification that each group exhibited. Agents on the edge of clusters have higher identification than those in the centre, this is a result of implementing the rejection-identification hypothesis. As agents on the edge of a cluster have more interactions with the out-group they have a greater chance of out-group rejection strengthening their identification with the group. This ties in to existing research suggesting that individuals in contact with their outgroup show higher identification and the finding that individuals seek to maximise the uniqueness of groups that they belong to (e.g. Brewer, 2003). A superordinate group provides less uniqueness as an individual encounters few people outside of it, whereas subgroups could have more frequent interactions with other subgroups leading to higher identification with these groups. DGM provides an account of how out-group rejection can facilitate this process, by causing individuals to identify more strongly with identities that make them distinct from others in their environment.

Network Propagation Outcome

Agents belonging to the group were initialised in a particular area. As demonstrated earlier there were only two conditions in which a larger number of group members settled in parts of the social network that were outside of this initial area. This outcome concerns how members of a group might propagate through a broad social network rather than having group

members exist in tight knit communities that have no effect upon individuals who are socially far removed. The concept of a social network discussed here describes a broad social network of many linked individuals and is not to be confused with the notion of an individual's social network, which describes the friends and acquaintances of a particular person. Two conditions enabled group agents to move to spaces on the grid outside of the initial area. One produced isolated clusters of group agents which existed outside of the initial area, but did not interact with the nongroup agents, the other lead to a more integrated population.

The first of these conditions required a very high level of positive ingroup interactions and a very high level of preference for similar agents. In this situation agents typically ended up forming 'enclaves' containing only agents belonging to the group and having no contact at all with agents outside of the group. The selection on these parameters in DGM is reminiscent the behaviour of cults and sects. The organisation of such groups enforces a high degree of social isolation and interactions within the community and this results in a lack of cross group contact and high identification with the group. In DGM this occurs under conditions involving a high preference for similar agents and a high level of positive in group interactions.

Members of organisations like cults and sects often weaken or eliminate existing social ties (e.g. Wright, 1991). Cults also exercise control over their members' social interactions (e.g. Wexler, 1995). DGM describes how these conditions lead to the observed behaviours of such organisations. However as the majority of social groups are not cult-like entities that force their members to have limited to no social interactions with other individuals this situation does not describe how most social groups might propagate through a social network.

The second condition may provide a much more typical indication of how a group propagates through a social network, in a manner that enables its members to interact with agents who are socially far removed. In this condition a low level of positive ingroup

interactions causes an individual who belongs to the group to seek social contact with those who are outside of it. A low level of preference for similar interaction partners makes this possible. The high level of contact between individuals who belong to the group and those who do not makes it possible for the group to recruit individuals from outside of its usual boundaries. It is worth noting that in these simulations, despite the group gaining recruits outside of its initial area, the group typically has fewer members overall. DGM would predict that the only way a group could both grow in terms of its absolute number of members and propagate through a social network would be under changing conditions, where the level of positive ingroup interactions varies over time in some fashion. Such conditions occur for real groups, or even within them as individual differences would cause some people to experience different levels of in-group interactions. It would be worthwhile for future research to examine these changes and try to determine how they impact on group formation.

Limitations and Outlook

The DGM uses a two dimensional grid to represent the social relationships between agents, which requires the conclusions drawn from the DGM to be interpreted with care, as they might not all hold up in a more realistic network architecture. However, the group propagation result aside, the results of the DGM are largely driven by the nature of the agents interactions rather than the topography of their social network. Additionally this representation provides a lot of advantages, in addition to practical concerns such as computational efficiency and ease of analysis it is a tried and tested approach used in a wide variety of existing successful models (see introduction).

The group simulated in the DGM is not representative of all groups, the assumptions used to drive the behaviour of the group limit the groups that it is possible to represent. For instance the group modelled acquires new members for conformity, an approach that would not make sense for some groups, for instance those based on gender or nationality. That being said the assumptions in the DGM represent the more common aspects of social groups so the model is still widely applicable. Taking distinct assumptions for the behaviour of the model is also necessary for the model to explain the target phenomena in terms of these assumptions. In this case group dynamics are explained in terms of social identity theory, the rejection-identification hypothesis and social conformity.

Conclusion

DGM demonstrates that social identity theory, the rejection-identification hypothesis and conformity taken together are able to explain a wide range of phenomena observed in social groups, from level of identification with superordinate groups to size limits of these groups. The sensible results it obtains in group coherency, integration and identification confirm that the simulation is reliable, allowing it to address its principle aim of testing the application of limited network sizes to social groups. In terms of obtaining a fixed group size, the DGM demonstrated how a social group's size is related to factors governing the interactions within the group, indicating that the line of research suggesting a fixed size for social networks could be applied to social group, reconciling sociological and biological theories of group size with current social psychological literature. How this applies in practice depends upon the group in question, as the social factors that govern interactions within a group are subject to change over a smaller time scale than average human neocortex size and

the social world is much more complex than is presented in this model. The next chapter embraces some of this complexity by expanding the model with further findings from the social psychology literature to better understand this relationship.

Chapter Four:

Extending Agent Based Models using Multiple Social Identities

Multiple Social Identities

The behaviour of the Dynamic Groups Model (DGM), which was introduced in the previous chapter, was based on social identity theory (Tajfel & Turner, 1979; Turner, 1987; Hogg and Abrams, 1998; Abrams and Hogg, 2000; Jetten & Postmes, 2006). In this model the behaviour of a single group was examined to answer questions about group dynamics. However social identity theory and related theories (e.g. social categorisation theory, Turner, 1987) hold that individuals simultaneously hold several social identities and that these identities are all relevant to how a group behaves. The impact of multiple social identities upon intergroup interaction has been an important area of study. For example Brewer, Ho, Lee and Miller (1987) explore the question of how people interact with individuals who share some, but not all, of their social identities. They used the medium of interpersonal judgements, exploring how predictions of ingroup bias map on to participants judgements of individuals who share some social groups. They find that individuals still express a bias favouring those who share their group identity, but the strength of this varies depending on the nature of the interaction. Combined with the work of Hewstone, Islam and Judd (1993), six different methods of combining information on group membership to form ingroup / outgroup judgements have been identified. Crisp, Ensari, Hewstone & Miller (2002) observed that in the absence of an experimental manipulation most participants used an additive method; an individual who shared groups A and B was evaluated more favourably than any other, one

who shared neither group was evaluated the least favourably and individuals who shared one of the two groups were somewhere in between. The DGM used a highly simplified version of social identity theory in many respects, but this aspect is particularly significant as all social research has necessarily been performed in the context of real individuals who hold many social identities.

The Current Research

This chapter focuses on expanding psychological simulations to account for multiple social identities. In general it is not advisable to expand such simulations with parameters that are not directly related to their findings, as this makes it harder to draw sound theoretical conclusions from the model (Smith and Conrey, 2007). Despite this, previous research has shown that if a parameter is implemented in the model as a matter of necessity then it is important to consider the consequences of varying that parameter (e.g. Gotts & Pollhill, 2010). The number of social groups in a simulation is such a parameter, as it is already implicitly implemented in most social simulations (including those presented in previous chapters) and as such the impact of it should be studied in more detail.

The first extension will be to expand the DGM presented in the previous chapter to focus upon multiple identities. This model had outcomes relating to several theoretical areas, but focused primarily on the relationship between the internal dynamics of a group and the stable size that the group reached. In the DGM only a single group was examined and agents either belonged to the group or did not. Modelling several groups and allowing agents to belong to some, all or none of them would produce more organic behaviour and generate

results that are closer to a true implementation of the underlying theory. The principle challenge for the extension, however, is not to see what new insights this might bring, but to examine whether the existing findings of the model change when agents hold multiple identities. This would support the notion that the multiple identities extension is necessary to ensure accurate simulations as opposed to simply being one of many possible extensions that might yield new insights.

After extending a model from this thesis the next step is to show that the multiple identity expansion has a wider applicability. This will be achieved by extending Schelling's (1971) classic model of segregation. This model demonstrated that large scale racial segregation can emerge from very mild preferences for similar neighbours. It has been important to the development of agent-based modelling as a field and is still being studied and expanded, for instance Henry, Pralat & Zhang (2011) recently applied the model to social networking. Schelling (1971) developed the model to explain segregation between black and white people in America, so it focused on just two social groups. Given that individuals do not define themselves as belonging to just one group, it is important to ascertain whether adding additional identities into the simulation might change the outcome. There is existing research which implies that the presence of cross-cutting groups reduces intergroup conflict. This was originally observed in sociology by Evans-Pritchard (1940), but the role that cross-cutting groups have in reducing prejudice have been widely replicated (see Brewer, 2000) and societies with more cross-cutting links experience less intersociety conflict (Fry, 2009). Extending Schellings (1971) model with additional social groups will explicitly model the processes that occur when such cross-cutting individuals are present.

Study One: Building on the Dynamic Groups Model

In chapter three the DGM had a social group which was compared to “nongroup” agents who did not belong to the group, these were assumed to belong to other groups that were not relevant for the purposes of the simulation. This extension to DGM will explicitly look at some of the groups that individuals outside of the studied group might belong to, however the concept of “nongroup” agents is retained, as countless social groups exist and those represented in the model will not be exhaustive. Three additional social groups are explicitly modelled, for a total of four modelled groups.

In line with multiple identity research (e.g. Crisp, Ensari, Hewstone & Miller, 2002) agents will be able to belong to more than one social group. This means that it is necessary to alter the process of joining groups described in chapter three so that agents can continue to join new groups even if they already have a group membership. To justify the introduction of multiple groups, the new model must fundamentally change the relationship between 'preference for like neighbours' 'degree of positive ingroup interactions' and 'final group size' identified in the previous chapter. If this occurs then it forms the basis of an argument for including multiple societies in more social agent based models.

Method

Extended Dynamic Groups Model (E-DGM)

As in chapter three, the model implements the interactions between agents from different social groups and tracks the progress of a social group in order to measure what impact the models parameters had. While four social groups are implemented, only one of them is tracked, this makes the results from the E-DGM directly comparable to the results from the DGM. This group is designated the “target group”. As the groups are implemented identically the results from the target group can be generalised to the other groups

As in chapter three, agents are placed onto 70% of the spaces on a 51 by 51 grid. Previously agents were divided into "group" and "nongroup" agents, now agents fall into one of three categories. “target group agents” belong to the target group, they may also belong to some or all of the other modelled groups. “group agents” do not belong to the target group, but do belong to some or all of the other modelled groups. “nongroup agents” do not belong to any of the modelled groups, but are still considered to have social identities on a theoretical level.

In the DGM a 20 by 20 area was populated by the group, in the E-DGM each of the four groups has their own 20 by 20 area. Half of the agents in a designated area will belong to that areas groups. The areas are placed equidistant from each other and so are each separated by ten rows of nongroup agents. This makes it possible to extend the DGM to examine multiple groups while maintaining an initialisation whereby the agents for each group are associated with a particular area. This allows for a comparison with the group proliferation

results in the previous chapter. The nongroup agents are now conceptualised as agents belonging only to social groups other than the four under examination.

In the DGM each cycle an agent that has too many neighbours who did not share its group would move, in the E-DGM this rule is adapted so that agents move when they have too many neighbours who do not share at least one group with them. This is still controlled by a parameter "Preference for similar neighbours". After moving, agents will interact with each other. An interaction with an outgroup member increases the agents identification with their group. An interaction with an ingroup member will increase the agents identification only if it is positive and identification decreases it if the interaction is negative (see chapter three for rationale). The chance of this occurring is controlled by the second parameter "Chance of positive ingroup interaction"

The interaction rules must be adapted to account for the possibility that agents belong to multiple groups. This is approached by having each agent in the interaction determine which of their identities is salient for that interaction, then the interaction proceeds as if each involved agent only had its selected identity. If two agents with the same salient identity are interacting then the interaction proceeds exactly as described for ingroup interactions in the previous chapter. If an agent with a salient identity is interacting with a nongroup agent this also proceeds as described in the previous chapter. If two agents with different salient identities are interacting then they both act as if they were having an interaction with a nongroup agent, as they are having an interaction with an agent that does not share their identity the same theoretical basis applies (rejection-identification). The agents select their salient identities in the following fashion: If they share at least one identity then they will both select the shared identity, if there is more than one shared identity then the identity that they both select is chosen at random. If they have no shared identities then they will select their

salient identities completely at random.

Following the interactions, agents may join groups that they do not belong to through conformity. If more than half of an agent's neighbours belong to the same group and the agent does not share that group there is a one percent chance per cycle that the agent will join that group, with an initially low identification. This rule is unchanged from the previous simulation, but now that there are multiple groups, conformity could cause a group agent to join a second group, developing multiple identities. The method by which agents might leave groups is also unchanged, sufficient ingroup rejection will cause an agent to lower its identification with the social group and ultimately leave.

In the previous chapter the simulation was completed after 1,000 cycles and the simulations were declared to be stable if less than 2% of the agents moved in the last ten cycles. However, in the E-DGM pilot studies reveal conditions in which the main outcome of the simulation (group size) would be stable while the individual agents positions would not be, which may have led to the misclassification of some simulations. In the E-DGM a simulation is said to have stabilised once the target group is not changing in size by more than five agents per hundred cycles⁸. It is possible that many more agents than this are changing identification, but the net effect upon the group as a whole is minimal.

Outcomes

⁸ As it is expensive to perform this calculation every hundred cycles in very long simulations, as the simulation went on the time between checks increased. For instance after 700 cycles the model would accept a group size variation of no more than 40 within the next 800 cycles as indicating a stable simulation.

As in the previous chapter the size of the target group is the primary outcome of each simulation, with the average strength of identification, level of segregation and number of agents outside of the group's initial area also measured to determine the validity of the simulation and draw wider conclusions. In addition to this the number of agents that develop multiple identities is also monitored, to assist in understanding how the model has been impacted by the inclusion of multiple social groups. The main hypothesis is that including multiple social groups will change the core findings of the DGM, with a focus on the finding that more positive ingroup interactions produce higher group sizes with the largest impact being at moderate levels of positive ingroup interactions. Additionally the E-DGM will examine the impact of multiple identities upon group dynamics that emerge from these simulations.

Simulations systematically examined the two parameters that drive the simulation ("Preference for similar neighbours" and "Chance of positive interactions") at values between 0 and 100 (inclusive) in ten point increments. Each combination of parameters was replicated ten times. In no case were agents with multiple identities were present at the start of a simulation.

Results

All simulations reached a stable final state. All of the outcomes were influenced by the agents capacity to form multiple identities, so the conditions which lead to agents having multiple identities are discussed first to contextualise the other results. Following this the

results from the DGM and the E-DGM are directly compared to demonstrate what influence the multiple identities extension has had upon this model.

Multiple Identities

Agents only formed multiple identities where the level of ingroup positive interactions was above 40. A much greater number of agents were able to form multiple identities where the preference for similar neighbours was high. The number of agents that exhibit multiple identities (defined as having an identification of greater than 0 in more than one group) can be seen in figure 4.1 below:




Figure 4.1: Agents with multiple identities

As no agents are initialised with multiple identities, these agents must emerge during the simulation. This occurs when an agent that already identifies with one of the four groups is adjacent to several agents from one of the groups it does not already identify with. There are two conditions necessary for this to occur, firstly the agent must not lose its initial group identity through ingroup rejection; secondly it must move to a position such that it can be outnumbered by members of the other group. The former does not occur if the level of ingroup positive interactions is low and the latter is more likely if the similarity preference is high, as this both motivates agents to move and causes the emergence of clusters of agents that would be dense enough to cause local conformity. The observed results reflect these behaviours.



There is an additional factor at work for simulations with a high level of positive interactions and a very high similarity preference. In these simulations thousands of agents develop multiple identities, in some cases all of the agents in a simulation belong to multiple groups. This occurs when the level of preference for similar neighbours is high enough that the agents cannot find a position that satisfies them, as such all of the agents move every cycle, occasionally finding positions that cause them to adopt additional identities. The process continues until there are a great number of agents with multiple identities which can then stop moving as they perceive more of the other agents to be similar to them, making it easier for them to satisfy their preference for similar neighbours. Taken to extremes, if all agents belong to all four groups then any level of similarity preference will be satisfied. The high level of positive interactions is necessary for this to occur, as it guarantees that the rate at which agents leave their groups through ingroup rejection is smaller than the rate at which new agents are joining groups through conformity. Where similarity preference and positive

interactions are both extremely high, every agent in the simulation eventually joins every group. As well as leading to more agents with multiple identities this process also leads to high group sizes and high numbers of agents settling outside of their initial areas.

The multiple identity formation results described above are helpful in understanding some of the reasons for the differences between the results of the DGM and the E-DGM. Other differences might also occur due to the mere presence of extra groups, any change in results where the 'positive ingroup interactions' are less than forty must be of this nature as no agents with multiple identities occurred in those simulations. The rest of the results for this study will be examined by directly comparing the results of the DGM to the E-DGM on each of the key measurements: the group's final size, the average identification score of members of the group and the number of agents outside the group's starting area.

Group Size

The key outcome from the DGM was that fixed final group sizes could be obtained and that they were determined by the degree of ingroup rejection and preference for like neighbours. The results for the final size of the target group in the DGM and the E-DGM are presented for comparison (See figure 4.2).






Figure 4.2: Group Size outcomes for DGM (top) full E-DGM (middle) and low similarity preference E-DGM (bottom)⁹


The in E-DGM have been able to reach a substantially larger size, sometimes managing to encompass of all of the agents within a particular simulation. For the low levels of preference for like neighbours (the darker lines) the overall pattern of the results is similar to the DGM; at low levels of positive ingroup interactions the group is small while more positive interactions lead to a larger group and the group size is more sensitive to changes in positive ingroup interactions around the 50 mark. Where high group sizes occur, a lower

⁹ When comparing these figures note the difference in scale.

preference for similar agents leads to a lower final group size, which is the reverse of the previous results. The reason for this is that a high level of positive interactions and high level of like neighbour preference create the conditions necessary for the process that leads to high levels of multiple identity agents described in the multiple identities section. When this process occurs a large number of agents convert to the group, regardless of their existing group membership, leading to the very large group sizes observed.

Identification

The introduction of the additional groups also changed the pattern of identification observed in the simulation. Figure 4.3 displays a comparison of the results from the previous chapter and the new results.






Figure 4.3: Identification Outcomes outcomes for DGM (top) and E-DGM (bottom)

As before, two distinct patterns of results can be observed, a linear positive relationship between the level of positive ingroup interactions and identification at a high level of positive interactions and a less clear relationship at a lower level. The results at the higher level of positive ingroup interactions is unchanged by the addition of multiple groups and can be explained in the same way, more positive ingroup interactions lead to more identification gains. While the group is sometimes much larger at these values the average identification of members is formed in the same way. A low preference for similar agents provides a slight advantage as it leads to a slight reduction in ingroup rejection, caused by an overall reduction in ingroup interactions. The remainder of this discussion will focus solely on the behaviour of E-DGM at positive interactions 60 and below, as this is where it deviates from the previous model.

In the DGM, a higher level of ingroup positive interactions consistently lead to a lower identification. This was due to ingroup rejection making clusters of group agents unsustainable, leading to the group agents being scattered and reinforcing their identities mostly through rejection-identification. In the E-DGM higher levels of ingroup positive interactions often produces higher identification results until a positive interactions value of 60 is reached.

The pattern of low identification that was obtained in DGM occurred because agents obtained a low identification through rejection, then became isolated and had no further interactions changing their identity (see chapter three for details). This does not occur in E-DGM because the agents never reach the point that they are isolated. When an agent was isolated in the DGM it could only be joined by another agent from its own group, which lead to more identity loss and one of them leaving the group until there were few enough agents that they could all find isolated spaces and stop moving. When an agent is isolated in E-DGM it is more likely to be joined by an agent from one of the other three groups. This increases its identification through outgroup rejection and maintains the same number of agents in the group, preventing there ever being few enough agents that they can all become isolated¹⁰.

See figure 4.4 for a demonstration of this process. Group agents are represented by red, green, blue and white squares, with brightness proportional to their level of identification. Nongroup agents are represented by dull grey squares. In this example no agent obtains multiple identities. Note that no pattern of stationary agents emerges and the group agents all develop high levels of identification.

10 At lower population densities it might be possible for all agents to become isolated despite remaining in the group, but the population density would have to be very low and high identification would still be obtained through rejection-identification as some of the agents interactions before it found the isolated location would be with the other three groups. In the DGM the low identification isolated pattern will always emerge as the group loses members until it has few enough to fit whatever space is available.




Figure 4.4: Similarity Preference 50 Positive Interactions 20 simulation at cycles 1, 10, 20, 30, 40 and 50 (top left to bottom right)

Proliferation

The frequency with which the group proliferates away from its starting position is greatly increased with the addition of multiple groups. In some cases this is simply an indication of constant movement, in which agents are essentially distributed randomly, but in many others a stable proliferation takes place. The influence of the parameters upon the number of agents found outside of the starting area is as follows:





Figure 4.5: Number of agents outside of the initial area for DGM (top) and E-DGM (bottom)

It is immediately clear that in even the modest level of proliferation observed at low levels of positive ingroup interactions in these simulations is greater than the highest level of proliferation observed in the single group simulation. This is caused by the effect discussed in the multiple identities section, whereby a great number of agents form multiple identities. Agents with multiple identities can easily occupy spaces that are not part of their group's initial area, as they can satisfy their preference for like neighbours using members of different groups that correspond to their other identities. Comparing figure 4.1 to figure 4.5 it is apparent that the conditions under which group agents proliferate through the social network are also the conditions under which there are a lot of agents with multiple identities.

Discussion

The key findings of the DGM in the previous chapter were that social identity theory, the rejection-identification hypothesis and conformity could be taken together to build a model of group dynamics that produced sensible results for group coherency, integration and identification while demonstrating that a fixed social group size can be obtained. The extension produced in this chapter builds on these findings, offering additional insight into process of social propagation and novel observations about the influence of multiple identities upon social group dynamics. Importantly the DGMs findings with respect to fixed group size are revised as under some conditions the findings of E-DGM do not support a stable group size.

Propagation

Compared to the DGM it is much more frequent in the E-DGM for group agents to reach an area outside of the one in which their group is originally dominant. Even disregarding the simulations in which the agents are constantly moving, and so cannot be said to be settled in any one area, the degree of propagation the group experiences is far in excess of the previous simulation. Agents with multiple identities are able to carry one identity to parts of the social space that it could not normally occupy, by using a second identity to maintain their new position. If the groups have low ingroup rejection this leads to agents converting each other to the identities that they spread. The importance of cross cutting groups in allowing a group to become accepted are well established (Brewer, 2000) but in describing conversion this result goes further, suggesting that it is beneficial for a social group to be compatible with as many other social groups as possible.

Multiple Identity Formation

Agents only obtained multiple identities in conditions where ingroup rejection was low. In this model all of the groups shared the same parameters, if one group had a high level of ingroup rejection all four did. The E-DGM did not simulate the consequences of groups with different levels of ingroup rejection, it is likely that under some mixed group conditions an individual with multiple identities could belong to at least one group with high ingroup rejection.

The result that multiple identities do not occur in conditions of high ingroup rejection

does have some important implications. The first is that the model produces more realistic results at lower levels of ingroup rejection, as this produces individuals that have multiple identities, a feature of the general population. Examinations of factors that behave very differently at high and low levels of ingroup rejection are best served by looking at the results where ingroup rejection is low. This is most important for the identification result, which behaves very differently at low and high levels of identification and suggests that the dominant pattern should be that identification strength linearly decreases with ingroup rejection. The second implication of the points at which multiple identification occurs is that groups with high ingroup rejection are incompatible with each other. Such groups would rarely share members, which implies that where such groups can be found they will be much less successful.

Group Size

The DGM's primary goal was to address whether social groups could have a fixed size, similar to that postulated for social networks by Hill and Dunbar (2002). It showed that under all conditions the group maintained a certain size. In the E-DGM this is still true, as the termination condition requiring the group not be substantially changing in size occurred in all simulations. However there is an important difference, in the previous chapter the group always stabilised at a size smaller than the number of agents in the simulation, but in the E-DGM some conditions lead to every agent in the simulation joining the social group. Under these conditions the model would predict infinite growth, simulating larger populations would lead to the additional agents adopting the group identity and forming a larger group.

This result is unrealistic, besides “human” there is no social identity that includes all individuals. This points towards the parameter settings that produce this result being values that do not occur in the normal course of events. As this pattern of infinite group spread occurred only where the preference for like neighbours was set to 90% or higher it can reasonably be dismissed as modelling a situation that does not occur in real social groups. Under more moderate conditions the E-DGM maintains the finding of the DGM that social groups are able to maintain fixed sizes for as long as the manner in which the groups members treat each other remains consistent.

Study Two: A Classic Model of Racial Segregation

Having shown that the multiple groups extension improves a model presented previously in this thesis, the next step is to apply it to the wider set of agent based models. This study replicates Schelling's (1971) model and extends it to include multiple social groups. This will achieve two things, the first is to test the hypothesis that extending agent-based models to examine multiple social identities will change the findings of the original model. The second is to study the consequences of introducing multiple identities for segregation, to test predictions in the cross-cutting literature (e.g. Brewer, 2000).

Existing extension

The Schelling (1971) model has previously been expanded to account for multiple social groups in Clark and Fossett (2007), however the goals of this paper diverge from those stated above and the analysis they provide is not sufficient to answer theoretical questions about the impact of multiple social groups on psychological agent based modelling. Their research is focused on simulating several racial groups and matching the simulations behaviour to survey data. This chapter focuses on implementing theories of social psychology, such as agents that have multiple identities and directly comparing the results from these to those of the original model.

Theoretical Background

Schelling's (1971) model served to explain why some American neighbourhoods were strongly segregated into black neighbourhoods and white neighbourhoods despite none of the residents having a preference for segregation. The model showed that for two social groups, a mild preference for being around similar individuals caused a high level of segregation. The multiple identity research introduced in the introduction to this chapter (Brewer, Ho, Lee & Miller, 1987; Hewstone, Islam & Judd, 1993 and Crisp, Ensari, Hewstone & Miller, 2002) will be used to expand this model. Existing research suggests that individuals who belong to multiple groups will reduce conflict between those groups (Evans-Pritchard, 1940; Brewer, 2000 and Fry, 2009), the extension will be able to contribute to this area of research by showing the impact of such individuals upon segregation (which is related to intergroup prejudice). The primary goal of the study is to test how effectively multiple identity extensions can improve existing modelling efforts.

Method

The Multiple Group Segregation Model (MGSM)

Replication. The base model was a replication of Schelling (1971), which used the following method: Agents from two groups are placed randomly onto a grid. Every cycle each agent chooses whether to move to a random empty position or to stay still. This decision is made by examining whether the proportion of adjacent agents that share its group is sufficient

for the agents preferences. Agents neighbour preferences are set globally (i.e. all agents have the same degree of preference) and are expressed as the minimum percentage of agents that must match the agent's social group to be acceptable.

Third Social Group. In order to examine the effects of more than two social groups the replication above is extended to add a third group. Instead of being assigned to one of two groups, the agents were assigned to belong to one of three groups. The implementation of the preference for like neighbours was unchanged, agents did not discriminate between the two groups that were different to them, their preference was still expressed as a requirement for a certain proportion of their neighbours belong to the same group as them.

Multiple Identities. To account for multiple identities some agents were assigned more than one identity. The number of agents assigned multiple identities and the types of multiple identities assigned are varied between simulations, but an agent would never be assigned less than one or more than two identities. Agents cannot gain or lose identities in this simulation so the number of agents with multiple identities remains constant throughout any given situation. When determining whether their 'preference for similar neighbours' was satisfied agents would treat any other agent that shared at least one group with them as belonging to the same group. As such an agent with two identities would only fail to satisfy their preference for similar neighbours if they were surrounded by agents that held only the third identity.

Terminology. As agents identities cannot be changed groups of agents will be referenced by which group(s) they belong to. The groups are arbitrary and are labelled the 'red' 'blue' and 'green' groups. Six types of agents are possible: red, blue, green, red-blue, red-green and blue-green.

Initialisation and Data Treatment

Each simulation was defined by three parameters: The preference for like neighbours, which determined the extent to which individuals wanted to interact with similar agents. The number of agents with multiple identities, which determined what proportion of agents in the simulation belonged to two social groups. Finally the identity balance parameter, which determined whether agents with multiple identities drew their identities evenly from all three groups or whether all agents with multiple identities identified with the same pair of groups.

The first set of simulations replicated the original study, to ensure that model was consistent with Schelling (1971). Only two groups were used in these simulations and the preference for like neighbours was varied between 0% and 100%. The second set of simulations tested the impact of extending the study to a three group study, all three social groups were used, agents with multiple identities drew their identities evenly from each of the groups. Again the preference for like neighbours was varied between 0% and 100%, but at each level the proportion of agents with multiple identities were also varied within the same range. The final set of simulations tested the impact of uneven multiple identities and was the same as the second set of simulations except that agents with multiple identities always had the same pair of identities (red-green). In all cases the simulation for each combination of parameters was repeated ten times.

A simulation was terminated when either all of the agents had ceased moving or after five hundred cycles. Any simulation that does not reach a stable state after five hundred cycles will be defined as not reaching an equilibrium. While technically every simulation will eventually reach a stable state, in practice this can take an extremely long time as some situations would require the agents to all spontaneously move into a perfectly segregated

position from a random state in a single cycle of random movement. Five hundred cycles represents a reasonable cut off point as almost all of the simulations that reached an equilibrium, did so within the first two hundred cycles. If a simulation reached a stable equilibrium, then the outcome from that simulation was the degree of segregation that had occurred. This was quantified by averaging all agents proportion of like neighbours.

Results

Replication

For the replication to be successful it must reproduce Schellings (1971) result that moderate levels of preference for like neighbours produce high levels of segregation. MGSM's results fit this pattern (See figure 4.6) with one exception: it does not produce a result if the like neighbour preference is very high. This is because the simulation reaches no stable equilibrium for 500 cycles. As described above, any simulation with a potential stable state (in this case 100% segregation) will eventually reach it, as there is a chance that random movements will naturally lead to it, but it would take a very long time. If run to conclusion the missing data likely represents very high levels of segregation, which is consistent with the overall pattern of this results and with the original Schelling (1971) study.




Figure 4.6: Replication of Schelling (1971)

Additional Social Group

Having confirmed that this simulation is able to replicate Schellings (1971) key results the next step is to study whether the inclusion of multiple identities produces any novel results, especially results which might have changed the conclusions drawn in the original study. Figure 4.7 shows the results of a simulation with three groups, but no multiple identities, compared with the results for the two group replication above. The results are similar, however smaller levels of segregation are obtained when the preference for similar neighbours is low. This occurs because when the preference for similar neighbours is low

enough the agents do not move and in a three group simulation the agents are initially less segregated. When the third group is added each agent will have, on average, $1/3$ of its neighbours belonging to the same group as it instead of $1/2$ as it was in the original study.




Figure 4.7: Extension of Schelling with an additional social group

Balanced Multiple Identities

Moving on to the results for the simulations in which agents had multiple identities; in most instances it is still the case that a higher preference for similar neighbours leads to a higher segregation and a modest preference can lead to high segregation (see figure 4.8 for a

comparison of results with different quantities of agents with multiple identities). The results are not identical, the inclusion of agents with multiple identities has had three effects upon the model.

Firstly, simulations with more multiple identity agents were able to find an equilibrium within 500 cycles at higher levels of similarity preference than simulations that had a low level of similarity preference. As more multiple identity agents are present it becomes easier to find a stable state in which all agents are satisfied, as agents with multiple identities can tolerate a wider variety of agents and can function as “borders” between areas of different types of agent. This causes agents with multiple identities to have a stabilising effect.

Secondly, the quantity of agents with multiple identities in the simulation has a strong influence upon how effectively a preference for like neighbours causes segregation. As more agents with multiple identities are present, the preference for similar agents has a smaller impact on segregation. If all agents have multiple identities then the preference for like neighbours has no significant effect on final segregation, $r(88) = -.02$, $p = .85$.

Finally, in general the addition of agents with multiple identities leads to a decrease in segregation. This is consistent with the long held finding that cross cutting categories increase tolerance (e.g. Brewer, 2000). The simulation obtains this result because agents with multiple identities can help members from groups that match both of their identities be content with their location, which makes them more likely to be in contact with each other. However, where the preference for similar neighbours is very low the pattern is reversed and higher levels of multiple identities lead to higher levels of segregation. In the simulation the cause of this is that multiple identity agents will use the identity that matches the agents around them when calculating the level of segregation. Note that this only occurs in groups that are already very tolerant (low preference for similar neighbours). The multiple group

segregation model would predict that in very tolerant groups introducing individuals who identify strongly with multiple relevant groups might increase the level of social segregation. The reason for this is that they will emphasise the groups that they all share, meaning that several secondary identities do not become relevant to the interaction reducing intergroup contact. However in the vast majority of cases and certainly in situations where there is even a moderate existing prejudice cross-cutting groups will reduce segregation.




Figure 4.8: Average segregation over all agents at different levels of multiple identities.

The three groups were implemented identically, so as one might expect, there are no significant differences between the patterns of segregation that emerge for each of the groups. Figure 4.9 shows the results for each of the three groups, they do not differ from each other or

the average of all agents above in figure 4.8.

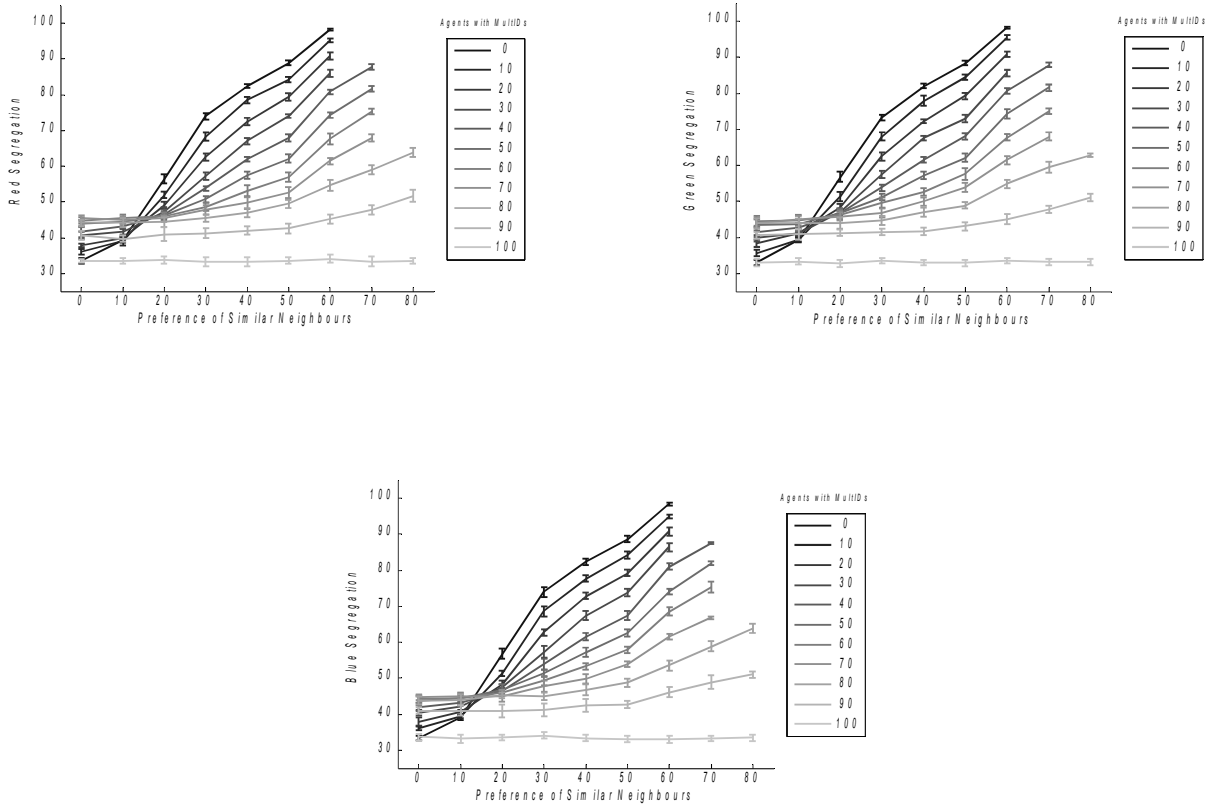


Figure 4.9: Segregation results demonstrating interchangeability of the red, green and blue groups.

Limited Multiple Identities

To better understand the impact of multiple identities, simulations were performed in which all multiple identity agents belonged to the red-green group. In these simulations no agents belonged to the red-blue or blue-green groups. This makes it possible to isolate the

effects of multiple identities upon groups which have more individuals with multiple identities as well as any influence they might have on groups that share no members with other nearby groups.

The most striking thing about these results is the absence of usable data at moderate levels of preference for similar neighbours that occurs when there are a lot of agents with multiple identities (See figure 4.11). These conditions lead to simulations that are unable to easily reach an equilibrium. The reason for this is that the agents from the blue group have trouble finding positions where their preferences are satisfied. As the red and green groups have the support of all of the agents with multiple identities they form large, combined red and green areas. Between certain points the preference for similar neighbours is too high to allow the blue agents to settle near these areas, but too low to force the areas to move away from the blue agents (which would create enough space for these agents to form their own areas). Unlike the previous simulations, the asymmetry between groups makes it possible for this simulation to reach a state from which it is impossible to form a stable state given any amount of time. Figure 4.10 below shows a simulation in which 40% of agents have multiple identities and agents require at least 30% of their neighbours to be similar to them. While some blue areas have stabilised, they are highly segregated and densely populated. This means that blue agents outside of those areas are still moving each cycle looking for a location that meets their preferences, but cannot find one as the existing blue areas are too densely populated to accommodate more agents and there is not enough space to form a new blue region. In order to confirm that the 500 cycle time limit of the simulation was not responsible for this result, the simulation was replicated with a longer time limit, it still does not reach equilibrium after more than thirty thousand cycles.




Figure 4.10: Multiple Ids 40 Preference 30 Cycle 32,715

Triangles represent an agent of their colour, yellow circles represent red/green agents

Aside from the absence of usable data at points where the simulation has not stabilised, the pattern of results is similar to the previous multiple identity simulation (see figure 4.8). The level of preference for similar neighbours after which additional multiple identity agents help rather than hinder desegregation, occurs at a slightly higher point. In the previous simulation this occurred at a preference for like neighbours of about 15%, here it occurs at the 20% mark.




Figure 4.11: The effects of agents with both red and green identities

In order to examine what impact the absence of blue agents with multiple identities has had on the groups, the relationship between similarity preference and segregation is plotted individually for each of the three groups (see figures 4.12 and 4.13). The red and green populations mirror the overall pattern observed in the simulation (compare figures 4.11 and 4.13) as was the case in the previous simulation. However the absence of agents with blue multiple identities has caused the blue group to differ substantially from the simulation average (compare figures 4.11 and 4.12). For the blue group higher preferences for similar neighbours still leads to higher levels of segregation, but the impact of the multiple identities parameter is different. Unlike the other groups adding additional multiple identities always leads to a decrease in segregation for the blue group, but this decrease is modest (see figure 4.12). The blue group also experiences much more extreme segregation scores than the other

two groups, varying between 0 and 100 as opposed to the 38 to 95 range occupied by the other groups. The blue group is only less segregated than the red and green groups where the preference for like neighbours is 20 or lower, which is the level at which an increase in agents with multiple identities leads to an increase in segregation for those groups.




Figure 4.12: Effects of red/green agents on blue segregation.

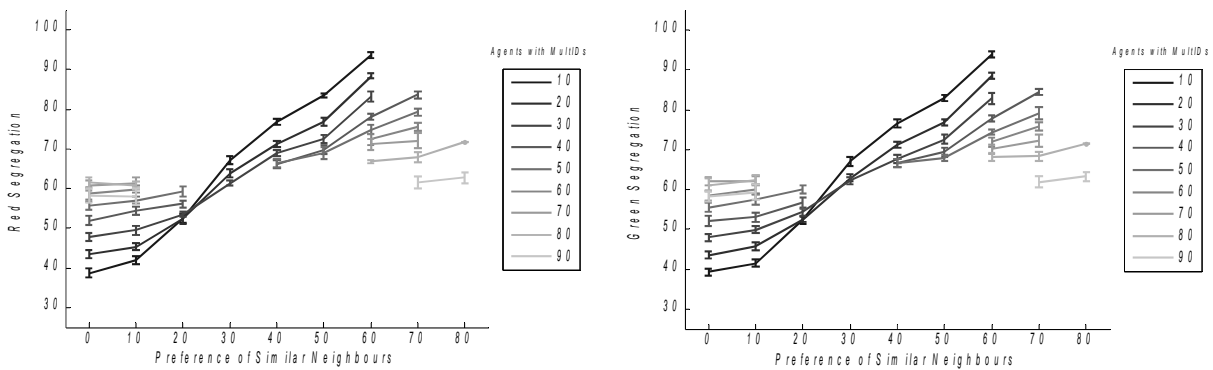


Figure 4.13: Effects of red/green agents on red and green segregation.

Discussion

Reinterpretation of Schelling

Extending Schellings (1971) classic model to account for multiple identities produced some novel findings and uncovers a situation in which the original findings of the model do not apply. The main finding of the original model is that a mild preference for similar neighbours can lead to a high degree of segregation, however this effect is weakened in the presence of three groups, as lower segregation scores are obtained. In a situation involving a large number of social groups, this effect would become smaller. This finding would suggest that the Schelling (1971) model may not accurately describe group processes where a very large number of groups are present. Most real world situations involve complex individuals with a great many social identities and it is obviously not the case that all social groups are completely segregated from one another. This suggests that Schellings (1971) segregation finding only applies in some conditions. Social identity researchers often find that only a subset of an individuals identities will influence a given interaction, specifically those that are made salient by the context of that interaction (e.g. Oakes, Turner & Haslam, 1991). Combined with the findings from this chapter this suggests that a more realistic interpretation of the segregation finding would be “Mild preferences for like neighbours lead to high levels of segregation, as long as only a few social identities are salient in the context of decisions relevant to segregation, such as housing decisions.”.

Effects of Cross-Cutting Groups

In addition to expanding on the Schellings (1971) main finding, the MGSM makes some novel predictions regarding the introduction of cross-cutting identities to segregated populations. Existing research (e.g. Brewer, 2000) supports the view that cross-cutting individuals reduce group conflict. However the MGSM demonstrates that under some conditions it is possible that this process might inadvertently marginalise a third group.

General Discussion

The conceptualisation of individuals as constructing multiple social identities for themselves has been highly influential in the development of social psychology (e.g. Turner 1987) and the implications of individuals with multiple social identities continues to be an area of interest (Crisp, Ensari, Hewstone & Miller, 2002). This chapter has expanded existing social agent-based models to take greater account of this facet of social psychology. Expansions of Schelling (1971) and the DGM presented in the previous chapter produce novel results. Extending Schellings (1971) model with multiple groups demonstrated the conditions under which the core finding of the model, that mild preferences lead to strong segregation, might not occur. This allows the model to explain why total segregation isn't observed in every case where a slight preference for ones in-group might be observed, a common preference (e.g. Tajfel & Turner, 1979). Instead the original finding is amended to specify that mild preferences will only create segregation when the group identities leading to the segregation are more salient than other relevant social identities.

Additionally this expansion results in a novel finding for the cross-cutting literature, which commonly holds that there are strong benefits to groups having individuals that cross-cut these groups (e.g. Brewer 2000), by highlighting a case where integrating two groups in this manner might inadvertently increase the isolation of a third group.

Extending the work of the previous chapter demonstrated conditions under which the fixed group size might not occur, albeit unlikely ones. It also drew attention to other important factors, namely the possible impact of other groups with different behaviours. Additionally, it demonstrated an incompatibility in being a member of multiple groups with high ingroup rejection. Finally it highlighted the important role of multiple identities in how a social group

obtains new members.

These findings expand upon the original studies and generate qualitatively different results, which provides strong evidence for giving more consideration to the impact of multiple social groups when creating agent based models. Considering more than the minimum number of groups and the possibility of agents belonging to multiple groups in previous and future modelling work will likely produce more reliable and more complete simulations.

Chapter Five:

Extending Evolutionary Simulations with Multiple Societies

The previous chapter presented two agent-based models that were extended to account for multiple social identities. The extended models produced novel results that, in some cases, directly contrasted with the original findings of the models. This demonstrated that applying principles from multiple social identity research to existing models can improve these models and their theoretical contributions. In the present chapter this approach is applied to addressing the paradox of altruism, the notion that altruism could not be an evolved trait as it does not contribute to the survival of the altruistic individual. Evolutionary simulations have been used to address this research area and the current chapter extends a recent model by Spronck and Berendsen (2009) with multiple social groups in order to explore the multilevel selection theory account of the evolution of altruism. As this is a different class of model to those presented in the previous chapter, an evolutionary simulation, the extension also tests how widely applicable the approach of studying additional social groups is.

Evolutionary Computation

Evolutionary computation is inspired by the theory of the evolution, which holds that biological organisms are adapted to their environment over generations through a process of natural selection. The theory is applied by implementing models in which agents interact in the context of some environment and are judged based upon how successful they are. Success

in the environment is termed 'fitness'; agents are considered more or less fit than each other and behaviours can be discussed in the context of how they impact fitness. Periodically, the least fit agents are removed and new agents are created by combining the traits of the fittest agents. Over a period of time this leads to the agents in the simulation becoming more adapted to be successful in this environment. The behaviours that lead these agents to become the most fit are termed adaptive behaviours. The goal of evolutionary models is often to determine which behaviours are most adaptive in a particular environment either to explain why these behaviours are observed or to determine the best approach to a particular problem. See Eiben and Smith (2003) for a general introduction to evolutionary computation.

Evolutionary computing has applications in a wide variety of fields from theoretical mathematics to machine learning. Biological simulations routinely show how organisms physiological characteristics are adapted to their environments, for instance Merilaita & Tullberg's (2005) simulation of camouflage vs costly signalling. Kameda and Nakanishi's (2002) model provides an example of how this approach is used in psychology. An evolutionary model is used to make predictions about human behaviour, in this case cultural learning, in order to advance theoretical knowledge of the domain.

In Kameda and Nakanishi's (2002) model agents would periodically perform one of two arbitrary behaviours, A or B. Each cycle their environment favoured one of these behaviours and an agent's fitness could be measured by how often it adopted the favoured behaviour. The initial favoured behaviour was random and the environment would occasionally change, such that the other behaviour became optimal. Each agent had "genes" that determined how it behaved, determining whether their choice of behaviour was more likely to be driven by their individual experience of the environment or social communications about the environment. An agent's genes also determined how likely the

agent was to contribute their experiences to other agents who were trying to learn about the environment socially. Kameda and Nakanishi (2002) found that when individual learning was expensive agents engaged in more social learning, but also that the quality of information gained through social learning was decreased. In this environment social learning was a more adaptive behaviour, even though in absolute terms it was less accurate. These results were then confirmed in an experimental setting where participants performed an analogous task in which they had to select which of two bushes a rabbit was hiding in. From this the researchers concluded that the conditions under which social comparisons were made are influenced by the nature of the environment, specifically focusing on the cost of making other types of comparison. Evolutionary models are useful in psychology because they follow this pattern of demonstrating that a behaviour is evolutionarily adaptive in order to predict and explain similar behaviours in human populations.

The Paradox of Altruism

The 'paradox of altruism' is a long-standing problem for evolutionary theorists. The challenge is to find an explanation for the persistence of altruism in human and animal populations. Evolutionary theory would predict that any behaviour that reduced the individuals chance of passing that behaviour down to the next generation (genetically or otherwise) would be removed from the population. However altruistic behaviours that reduce the individual's fitness in order to benefit another's are relatively common.

Several explanations for the emergence of altruism have been proposed, examples include reciprocal altruism and kin selection. Reciprocal altruism describes altruism which

persists because altruistic agents benefit each others chance of survival but contribute nothing to outsiders (Trivers, 1971). Kin selection describes a process in which altruism is genetically beneficial because while the individual suffers it offers benefits to genetic relatives (Hamilton, 1964). See Sober and Wilson (1999) for a general introduction to theories of evolutionary altruism. Warneken and Tomasello (2009) review the psychological evidence surrounding the emergence of altruism in human children and in primates. They conclude that children start to develop altruistic behaviours at a younger age than they develop behaviours that might support the emergence of altruism, such as forming and evaluating reputations (A process necessary for reciprocal altruism). This evidence does not disprove theories based on individuals undertaking actions that support other altruists, such as reciprocal altruism (Trivers, 1971) but it does challenge them to offer an explanation of how individuals come to undertake altruistic behaviour before they are capable of performing the actions required to maintain altruistic behaviour in the population. Given the time-scale involved it is not possible to experimentally study the evolution of altruism in humans, so it is difficult to integrate these explanations, but evolutionary computation approaches make it possible to investigate this area.

This chapter extends Spronck and Berendsens (2009) model of evolutionary altruism. Spronck and Berendsen (2009) examine the effects of memory, kin recognition and reputation using an agent-based evolutionary model to simulate a society. In their model agents must gather food to survive, which they can do by foraging for it, stealing it from other agents or receiving it as gifts from other agents. The goal is to understand under what conditions agents will engage in altruistic behaviour (i.e. giving food away). In one condition agents are able to remember specific previous interactions and use this to drive their behaviours, in another agents are aware of each others general patterns of behaviour, either of these could produce

altruism in the manner predicted by reciprocal altruism. In a third condition agents can distinguish which other agents are their kin and use this to drive their behaviours, which could produce altruism as predicted by kin selection. While Spronck and Berendsen (2009) found that most of these conditions lead to a reduction in 'destructive egoism' (agents stealing food from each other) only one condition lead to an emergence of altruism (agents giving food to each other). Altruism emerged in the 'friendly forever' reputation condition, in which individuals were aware of each others reputation, which was positive as long as the agent had performed a single altruistic act, regardless of their other behaviour. Spronck and Berendsen (2009) also explicitly state that they are examining a single society, which makes their model a good candidate for replication and extension here, as extending the number of societies investigated allows for the theme of examining multiple social groups to be continued from the previous chapter. Additionally their mixed results mean that the extension can be evaluated both in terms of its impact on the existing positive results (altruism emerging in the friendly forever condition, general decreases in destructive egoism) and whether novel results are produced in other areas (such as altruism emerging in the other reputation conditions).

A theory of the emergence of altruism not addressed in Spronck and Berendsen (2009), but that is relevant to a multiple society extension, is multilevel selection theory (e.g. Wilson & Sober 1994). This holds that members of one group may be selected over another if the group as a whole is more evolutionary fit, even if its individual members are not. More precisely, altruistic behaviours increase fitness because the ability to belong to a beneficial group contributes to an individual's evolutionary fitness. The consequence of this is that individual behaviours that are nominally bad for the individual, but good for their society might be preserved by virtue of the success of societies containing these individuals. There is support for this notion in the agent-based modelling literature. For instance Nahum, Harding

and Kerr (2011) demonstrated that multiple societies of competing agents with relevant differences between them can produce altruistic behaviours. In this model agents are in a “rock-paper-scissors” relationship, in which each type of agent has an advantage over one other type of agent, but is disadvantaged when competing with the third type. Nahum, Harding and Kerr (2011) simulate a group of toxin producing bacteria that are able to successfully invade a population of toxin sensitive bacteria, but be easily invaded by a population of toxin resistant bacteria. This model produces altruism, which in this case occurs when a bacteria restricts its own growth, in a way that is consistent with observed natural phenomenon. This demonstrates that multilevel selection theory has implications for the emergence of altruistic acts in simple entities, the next step is to apply the theory to a simulation of human development.

The Present Research

The simulation presented in this chapter extends Spronck and Berendsen's (2009) model of the evolution of human altruism to explicitly include multiple groups. This serves two purposes, the first is to test whether multilevel selection theory can lead to more altruistic patterns of behaviour in a simulation of human behaviour. The second is to contribute to the theme started in the previous chapter: Extending agent based models to consider multiple social groups. The remainder of this chapter will focus upon describing a multi-society extension to Spronck and Berendsen's (2009) model, simulating the effects of intersociety competition on the emergence of altruism and discussing these results in the context of the altruism and agent-based modelling literatures.

Method

Baseline Simulation

Agents. The original simulation, described in Spronck & Berendsen (2009), was implemented as follows. Initially 300 agents are distributed randomly on a 30x30 grid. Each space on the grid is characterised by its location and the quantity of food present. There is no limit to the amount of food that a space can contain. Each agent is characterised by its age, its health and its internal rules. An agent's age is measured by counting the number of cycles that it has survived and ranges from 0 to 100, an agent dies of old age after 100 cycles. An agent's health is a measure of how much food it had eaten minus the amount of energy it had used to take actions, there was no limit to how much health an agent could obtain. An agent's internal rules determine how the agent behaves, a process described in detail below.

Actions. In each cycle an agent takes one of five actions: wander, forage, steal, share or mate. A newly born agent does not act on the cycle on which it is born. Agents can only see other entities within three orthogonal squares; the area that an agent can see is a roughly circular area centred on itself consisting of twenty-five grid spaces. A wandering agent moves to a random unoccupied tile that it can see. A foraging agent moves to random unoccupied tile that it can see contains one or more food. A stealing agent lowers another agents health by 25 and increases its own health by the same amount, if the target agent has less than 25 health the stealing agent gains less than 25 health. A sharing agent lowers its own health by 25% of its current total and increases another agent's health by the same amount.

Agents take the mate actions in pairs. A new agent is created with an age and health of 0 and is placed in a random space adjacent to one of the parents. Both parents immediately transfer 25% of their current health to the new agent, as if they had taken the share action, though this mandatory share action is not counted for the purposes of determining if altruism emerges in the simulation. The internal rules for the new agent are determined by combining the rules of both parents, described in the combining rules section below.

Internal Rules. Each agent has a set of internal rules that determine how they behave. Agents have 5 to 8 rules, from a set of more than 150 million possible rules¹¹. Agents are not able to change their rules, new rules will only emerge when a new agent is born. Each rule consists of a precondition and an action, when a rule's preconditions are met, the agent will take the associated action. If more than one rule has its preconditions met in the same cycle, the agent will follow the rule with the more restrictive precondition. The action associated with a rule is stored in a single gene, a precondition consists of 0-2 logical statements, each of which requires four genes. An example rule that has a precondition containing one statement might be:

“self health less than 25 steal”¹²

The first four genes describe a logical statement, the fifth and final gene determines which action the agent will take if this logical statement is true. The statement's genes determine who an agent observes, which feature of that entity they observe, a value to

11 Many of these rules are functionally similar, but even taking this into account results in thousands of distinct rules.

12 A reader familiar with genetic algorithms will recognise this as a simplification, this rule would actually be expressed as [1 1 2 3 25 1], but as inheritance is handled using the uniform crossover operator which works just as well with nominal values as it does with ratio values the simplification does not detract from the explanation.

compare that observation to and the type of comparison made. An agent can observe itself, the nearest agent, the furthest agent, the agent with the most health or the agent with the least health (selected from visible agents i.e. those within three orthogonal squares). In the example above the first gene indicates that the agent will observe something about itself.

The second gene determines what the agent will observe. The agent will observe the age or health of the agent selected by the first gene, or ignore the target agent (making the first gene irrelevant) and identify how much food is on the space with the most food (selected from visible spaces). The example rule above calls for the agent to observe the health of the selected agent.

The third and fourth genes determine a constant that the observed value will be compared to and the type of comparison that will be made. The constant takes an integer value from 0 to 100 (inclusive) and the comparison will be 'equal to' 'greater than' 'less than' or 'wildcard'. The third gene of the example rule means that the agent will check if its health is less than some value and the fourth gene sets that value to be twenty five.

The final gene determines whether the agent will steal, share, wander or forage if the precondition is fulfilled. The example rule calls for the agent to steal. Taken as a whole, each cycle the agent will check its own health and if it has less than twenty-five health remaining it will steal from a nearby agent. In other words the agent will steal when it is running out of food.

When a precondition has two logical statements, the precondition is not considered to be satisfied unless both statements are true. If a precondition has no logical statements it is always considered to be satisfied. Such a rule would just contain an action, that the agent would consider taking every cycle. However, as mentioned earlier, if the preconditions for more than one rule are satisfied, the agent will always select the action associated with the

rule that has the most restrictive preconditions. A precondition with more statements is always considered to be more restrictive than a precondition with less statements, thus a rule with no preconditions will only be followed if the other rules are all rejected or also have no preconditions. If two preconditions that have the same number of statements are fulfilled, then the operator called for in the third gene of the statements is used to determine which rule will be followed. The 'equal to' operator is considered more restrictive than the 'less than' or 'greater than' operators, which are in turn more restrictive than the 'wildcard' operator. The wildcard operator ignores the observed value and the constant, instead it always evaluates the comparison as true. However a rule with a wild card could still fail if the observation could not be made in the first place, for instance an attempt to observe the nearest agent where no other agents are close enough for the acting agent to see. Consider the following example of a complete set of internal rules:

- 1 “self age equal 10 strongest health less than 25 steal”
- 2 “self health greater than 100 weakest health less than 10 share”
- 3 “self age greater than 90 share”
- 4 “furthest food wildcard 10 forage”
- 5 “wander”

If this agent were twenty cycles old, near to a patch of food and one other agent was nearby with twelve health, then the preconditions for rules four and five would be evaluated as true. This will lead the agent to forage on the basis that rule four is more specific than rule five, as it has a precondition with one more statement than rule five. If the agent were instead

ten cycles old then rule one would also be evaluated as true and the agent would steal, as rule one is the most specific rule, having two preconditions and using the equality operator. If, instead, the other agent were not present then the agent would wander, as the precondition for rule four would fail when observing the furthest agent is not possible. This occurs even though the rule would ignore the designated agent and look for food on nearby tiles, allowing the agents to evolve rules that change their food searching behaviour based on the presence of absence of other nearby agents. If the first gene of rule four had called for an observation of the self then the behaviour would not change based on the absence of a nearby agent.

In the event that the agent is following a rule that leads to an action requiring a target (steal or share) then the agent will preferentially target any agent that was observed to satisfy the precondition. If no such agent exists, or if the agent observed itself then it selects a target at random. If multiple agents are selected this way, for instance due to observing two different agents to fulfil a two statement precondition, then the agent selects one at random. In the example above, if the agent were ten years old and the strongest agent had a health of less than twenty-five then the agent would steal from the strongest agent, regardless of any other agents that might be present.

Universal Rules. In addition to their internal rules, there are two rules that are universal to all agents. The first was that if a pair of agents have an age of at least 18, a health of at least 50 and can see each other then they will both ignore their internal rules and instead take the mate action. The second is that if the preconditions of all of an agents internal rules are false or the agent has selected an impossible action (for instance stealing when no other agents are present) then they will take the forage action.

Combining Rules. When a new agent is created by two other agents mating, the new agent generates its internal rules by combining its parents rules. In this instance each gene in

the new set of rules would randomly copy one of the parents, with equal probability. For instance if an agents parents first rules were "Self health greater than 50 share" and "Nearest agent health less than 25 steal" the new agent might adopt a rule such as "Self health less than 50 steal". Each inherited gene had a small (2% chance) chance of mutation, in which case it was replaced with a random valid gene instead of being inherited from either parent.

Food Growth. At the start of each cycle, before agents take actions, some new food is added to the environment. The proportion of spaces selected to produce food changed from cycle to cycle, the proportion of spaces producing food followed a normal distribution with a mean of 0.06 and a standard deviation of 0.001. Each space that produced food, produced twenty units of food, which were added to any food already present on the space.

Eating and Ageing. At the end of each cycle, after the agents have take their actions, every agent that is on a tile with food consumes some of it. In one cycle an agent consumes twenty units of the food on its space, or as much as is available if the space has less than twenty units of food remaining. The agent gains one health for each unit of food consumed. Then each agent loses one unit of health, dying of starvation if it reaches zero. Finally, each agent adds one to its age, dying of old age if it has reached one hundred.

Initialisation. At the start of the simulation one percent of tiles contained one hundred units of food, the remaining tiles contained no food. 300 agents were randomly placed onto the grid, no two agents could occupy the same position. Agents started with an age of 0 and a health of 100. The agents initial rules were determined randomly.

Data Treatment. Each simulation runs for 20,000 cycles and the actions that the agents take in the last 5,000 cycles were recorded. The frequency with which the agents took the four main actions (mate is ignored as it is never called for by the agent's internal rules) is used to measure which behaviours emerged over the course of the simulation. This is a more

accurate measure than examining how often the behaviours occurred in the agents' internal rules, as some preconditions will be fulfilled much more frequently than others. Altruism was said to have emerged if agents selected the share action more than one percent of the time. Some definitions of altruism would extend to considering the forage action as altruistic, as it produces less food for the agent than stealing and serves to improve the overall resources of the society. However Spronck and Berendsen (2009) adopted a stricter definition by which only the share action qualifies, as the forage action does not directly reduce the fitness of the agent performing it. In order to remain consistent with the original simulation the same definition is adopted for the multi-society extension.

Reputation Condition.

Spronck and Berendsen (2009) expanded their baseline simulation to examine the effects of other factors, such as altering the preconditions to make it possible for agents to recognise biological kin or to remember past interactions. In this chapter only the reputation conditions are replicated, as one of the reputation conditions was only condition to successfully produce altruistic behaviour in the original simulation (agents selected the share action in at least one percent of cases). In the reputation conditions agents gained a reputation based on their sharing or stealing behaviour. In the basic reputation condition an agent's reputation is 'friendly' if it shares more than it steals, 'neutral' if it performs both actions evenly and 'hostile' if it steals more often than it shares. Two additional reputation conditions were investigated. In one an agent is marked as 'hostile' forever, after a single incidence of stealing. In the other an agent is marked as 'friendly' forever, after a single incidence of sharing. Spronck & Berendsen (2009) found that only the friendly-forever condition gave rise

to significant levels of sharing. Under the reputation condition the structure of the agent's internal rules is changed such that instead of the first gene of a precondition's statement representing 'self' 'nearest/furthest agent' or 'strongest/weakest agent' they would instead select 'self' or 'agent with positive/neutral/negative reputation'. For instance an agent could develop a rule such as “Positive reputation health less than ten share” which would cause it observe a nearby agent with a friendly reputation and share with it if it were running out of health.

Inter-society Competition Model (ICEM)

The model described above is extended into a model of inter-society competition, the intersociety competitive evolution model (ICEM). The simulation described in Spronck and Berendsen (2009) is run four times simultaneously, to represent four different regions, each containing a different society. Agents from one society do not have one to one interactions with agents from any other society, instead inter-society interactions are simplified by modelling the effects of competition between societies. This interaction represents periodic struggles for resources at the society level, such as villages competing for control of the best hunting grounds, or nations fighting over their borders. Societies that are more fit as a whole win these conflicts, which provides a benefit to the members of that society. The benefit is applied to all members of the society, but as agents are able to evolve to steal or gift these rewards, other social structures may emerge over the course of a simulation. In the original model the fitness of an individual agent is measured by its health, so in ICEM the fitness of a society is measured by the total health of all of its agents. At the end of each cycle the societies interact, with the most fit society winning this competition and the least fit society losing it. All of the agents in the winning society gain one health, representing the resources

newly acquired by that society. All of the agents in the losing society lose one health, representing the resources that they have lost.

Initially simulations are performed with no inter-society competition, in order to confirm that ICEM produces similar results to the original model. In line with Spronck & Berendsen (2009) a set of ten simulations was performed for each of the following conditions "No reputation" "Basic Reputation" "Friendly Forever" and "Hostile Forever". The frequency of each behaviour (wander, forage, steal and share) is recorded in the last 5,000 cycles. A significant level of altruism is said to have emerged if agents choose the "share" behaviour more than one percent of the time. This threshold is required as the two percent mutation rate might lead to small amounts of altruism even if altruistic agents typically die out before reaching mating age. In order to be a successful replication, in the absence of competition each society in ICEM should produce altruistic behaviour under the "Friendly forever" condition, but not under the other three conditions.

Following this, further simulations test the effects of competition. Again simulations are performed for the "No Reputation" "Basic Reputation" "Friendly Forever" and "Hostile Forever" conditions and each condition is replicated ten times. However in this case the results are organised differently, the frequency of wandering, foraging, stealing and sharing behaviours over the last 5,000 cycles are observed individually for each society. The frequency with which the different societies win the intersociety competition over the course of the simulation is also recorded. This makes it possible to examine whether any emerging altruism is related to the intersociety competition, by comparing the level of altruism in societies that have won many competitions to those that have lost frequently. Should such relationships emerge, investigating the patterns of behaviour that occur in the simulations will clarify why altruism is able to emerge under these conditions. For instance whether altruistic

agents are more likely to occur in societies that perform poorly or well in the intersociety competitions, or the casual direction of any relationship between the emergence of altruistic agents and performance in inter-society competitions.

Results

General Pattern of Evolution

When the simulation is initialised each of the four target behaviours (wander, forage, steal and share) are equally likely to occur, as the agents internal rules are random. However patterns quickly emerge, in which some behaviours become much more common than others. This is because some behaviours are more adaptive than others; adaptive behaviours make the agents exhibiting them more likely to reach mating age with enough health to mate and pass on some of their rules to the next generation, agents exhibiting maladaptive behaviours do not, so the adaptive rules are more common in each consecutive generation.

For example, an agent's internal rules might favour a specific action, due to having more rules that lead to that action, or a rule with preconditions like “Self age greater than zero and self age less than one hundred” which is practically guaranteed to be satisfied and is highly specific (as the precondition has two statements). Rules that lead to actions that gain the agents health (steal or forage) are adaptive. The agents that have these rules will be more likely to have more than fifty health when they reach age eighteen and be able to pass on the rules to the next generation, making them more common. Conversely rules favouring actions that do not gain the agent health (wander) or cost them health (share) do not lead to the agents

having enough health to mate when they become old enough to pass on their rules. Over the first few hundred cycles of the simulation many of the agents who favour wander and share actions are eliminated without mating, while those favouring forage and steal mate often, causing these rules to become common. Rules calling for the adaptive forage and steal actions (such as “if nearest health greater than zero steal”) become much more common than rules calling for wandering or sharing (such as “if self age greater than five share”).

As the simulation goes on the effectiveness of the agents increases, as more and more adaptive rules are passed on to the next generation where maladaptive rules are lost. This leads to the agents facing stiffer competition for limited resources, as the same amount of food is added to the simulation each cycle. Note that the reverse is also true, if the random initialisation of the simulation happens to consist primarily of agents with maladaptive rules many of them will die off. This leaves a relatively massive amount of food for the remaining agents who will almost certainly survive despite their maladaptive behaviour. Eventually random mutations will lead to agents with more adaptive rules being born and the process described here will still happen, albeit somewhat later in the simulation. Regardless of the initial state of the simulation, at some point agents will improve to the point that favouring the actions that acquire more food is common enough that more complicated sets of behaviours are required to flourish in the environment.

An agent with the zero statement precondition rule “steal” might do well in the early stages of the simulation, as it will sometimes steal, gaining twenty-five health in a cycle. An agent with the rule “If nearest agent health greater than twenty-five steal” might also do well, as while it steals less frequently, when it does it guarantees that it will obtain the maximum amount of food possible from a steal action (agents cannot obtain more health in a steal action than their target had). However an agent with both of these rules would be even more

effective, as they would preferentially steal from the nearest agent if that resulted in an optimal steal action, but would still steal from a random nearby agent if that were not possible. Agents that have the most adaptive sets of internal rules, made up of rules with useful preconditions (as opposed to ones that are simply easily satisfied) obtain more of the limited food and mate more often, passing on their rules more frequently than other agents. As the simulation continues this leads to these more adaptive rules becoming increasingly common; over the course of the simulations agents are able to develop highly complex and well adapted internal rules.

With around 100 agents having 5-8 rules of 1-9 genes, it is impractical to manually track the development of the agents rules over 20,000 cycles. Examining a particular agent in detail would also be unhelpful as any individual agent would die after at most one hundred cycles. Instead it is useful to look at the patterns of behaviour that emerge, as in the original Spronck and Berendsen (2009) study. They used the frequency with which agents execute the four actions as a measure of which actions are called for by the most adaptive rules. Additionally examining the types of rules that emerge helps to understand how these patterns of behaviour become common. The large set of possible rules means that it is uncommon for agents to have exactly the same rules, but often surviving agents have rules that are functionally similar, for instance “if self health less than 25 steal” behaves similarly to “if self health less than 23 steal”.

Important Emergent Behaviours

Interpreting the analysis that follows requires an understanding of two behaviours that emerged in some of the simulations, to facilitate the discussion of the conditions in which they emerged. The first of these is that the population of some of the societies was unstable, which had implications for that society in the intersociety competitions. The second is the emergence of a series of rules that had the effect of punishing agents with negative reputations, this is critical to the emergence of altruism in this model.

The population size of each of the regions in each simulation never remained stable, every cycle some new agents are born and some old agents die so the size of the population changes over time. The magnitude of these changes is something that varies between simulations and between regions within the same simulation. Some populations will remain fairly close to their average size, others will experience large population surges and crashes. Figure 5.1 shows the population size over time for each of the four regions in a sample simulation¹³, each region is inhabited by a single society. The populations of regions 1 and 2 undergo more change than regions 3 and 4. As population size is closely related to total health, this is causing the less stable populations to both win and lose more of the intersociety competitions than the other regions, which affects the behaviours that agents in those region develop. This provides an explanation of why the correlation between how often a society wins and loses the intersociety competitions is sometimes not as highly negative as might be predicted from the facts that a society cannot both win and lose the same competition and that winning leads to an increased chance of winning in the subsequent cycle (as agents in the society gain health).

¹³ The sample simulation uses the basic reputation condition with intersociety competition, but the traits observed are common to simulations that do not share these parameter settings.




Figure 5.1: Sample population over time for 500 cycles (basic reputation condition)

Another key property of the simulations is the emergence of punishing agents. These agents have rules which cause them to steal from agents with a negative reputation, the prototypical example would be “negative reputation (any observation) wildcard (any constant) steal” which would cause the agent to identify a nearby agent with a negative reputation and steal from it (The variable and constant are irrelevant due to the wildcard). There are several rules that fit this category, for instance they could be combined with other adaptive preconditions. An example of this would be “negative reputation health greater than twenty-five steal” which would cause the punishing agent to steal from another agent with a negative reputation, but only if the punishing agent would gain the maximum possible health from the

steal action. For the purpose of this analysis a rule is designated as creating a punishing agent when its precondition contains a statement with the target gene “negative reputation” and the action gene “steal”. Such rules would cause an agent to identify another agent with a low reputation and possibly steal from them, depending on whether the other requirements of its preconditions were met. The frequency with which these rules occurred in the basic reputation and friendly-forever simulations is presented below (see Table 5.1). Stealing behaviour is common in all simulations, the majority of stealing rules are not influenced by reputation. Where agents do use another agent's reputation to determine who to steal from, the majority will target hostile rather than friendly agents.

Table 5.1: Frequency of different types of stealing rules in the final cycle of the simulation.

	Total number of rules	Frequency of stealing rules	Frequency of rules stealing from hostile agents (punishing rules)	Frequency of rules stealing from friendly agents
Ten "Basic reputation" simulations	33551	15003	4538	1166
Ten "Friendly forever" simulations	31220	8735	4950	652

The role of agents that enforce social norms by punishing transgressors has been established in other agent based models, for example Ord (2006) studies the role of enforcement agents in the iterated prisoners dilemma (see general discussion for details). ICEM demonstrates how punishing behaviours become commonplace, through the following process: Punishing agents are able to emerge and thrive because they are likely to gain more

health and thus be more successful overall than agents who steal indiscriminately. If an agent could steal from two other agents, one of which is likely to perform the share action, it is better for the agent to steal from the agent that will not share, as the stealing agent may receive food from both agents. A sharing agent gives up 25% of its current health, stealing from the sharing agent will reduce the sharing agents current health, causing the stealing agent to receive less health from the share action. Where a choice is available stealing from low reputation agents who are less likely to share is adaptive. If no agents share then all agents will have a negative reputation, in which case a punishing agent is functionally identical to an agent that steals indiscriminately. The combination of punishing behaviour being adaptive when sharing agents are present and at no disadvantage when they are not allows punishing behaviour to emerge even when very few agents are sharing. Once punishing agents are common, a positive reputation contributes more to an agents overall fitness as it leads to being stolen from less frequently. This causes sharing to become a more adaptive behaviour as it leads to a positive reputation. When sharing agents can survive long enough to mate this produces a systematic rise in sharing behaviour. These processes involve advantages for the agents adopting the punishing and sharing behaviours and explain how punishing and sharing behaviours co-evolve. Sharing behaviours do not necessarily provide a society level advantage, in some instances it can hinder a societies ability to function in the intersociety competitions. As these competitions are based on the total health of the society sometimes foraging agents contribute more than sharing agents as foraging allows the society to acquire new resources rather than redistributing existing resources. The impact that emerging punishing and sharing agents have on their societies is discussed in detail alongside the simulations in which they emerge.

Replication

With the intersociety competition disabled, ICEM reproduces the key results of Spronck and Berendsen (2009) (see table 5.2). The key finding that the sharing behaviour only occurs more than one percent of the time under the “friendly forever” condition. While these results are not identical to the original study they replicate its most important points, thus the conclusions drawn from the original model stem from the assumptions that underlie it rather than a specific implementation¹⁴. These results show that ICEM replicates Spronck and Berendsen (2009) correctly and provides a baseline to which the intersociety competition extension can be compared.

Table 5.2: Results of Altruism Replication (percentage of all actions represented by each type)

	No Reputation	Basic Reputation	Friendly Forever	Hostile Forever
Wander	0.85	0.95	0.34	0.94
Forage	59.50	59.73	70.71	59.42
Steal	39.18	38.63	21.26	39.21
Share	0.47	0.69	7.7	0.42

Spronck & Berendsen (2009) did not perform a detailed analysis of their agent's

¹⁴ There are always small differences between a replication of a simulation and the original, where the author of the original has not specified some technical detail. For example two rules might be satisfied that are as specific as each other, one implementation might use the first rule in the list and the other implementation might randomise the selected rule. In the former case agents would adapt to have more important rules occur sooner in their ruleset, but this adaptation would be absent in the other simulation producing slightly different results.

The important thing while interpreting these models is that the observed target behaviour is a result of the implemented environment. If the conclusions drawn from a model are accurate then the key results should occur.

when the relevant theories are implemented, even if small changes exist in the exact nature of the implementation. This is exactly what occurs here. The fact that we observe the same patterns in the data merely shows that these conclusions are not an artefact of how these small details are implemented.

internal rules to determine why altruistic agents emerged, but it is likely that the punishing agents that emerged in this replication were also present in the original simulations and that this facilitated altruistic behaviour. When having a positive reputation has a health benefit (some agents do not steal from the high reputation agent) the health loss of altruism (giving away some food) is mitigated to the point that a few agents adopt some altruistic behaviour. As the cost of establishing a positive reputation is much lower in the “friendly forever” condition (an agent only needs to give food away once) this is where the highest level of altruism is observed.

Intersociety Competition: No Reputation

The first test of the inter-society competition condition was to expand the no-reputation baseline model to examine four societies. Each simulation was repeated ten times, each one reported the results for four societies. The results for the societies are considered separately, as the inter-society competition manipulation may cause one society to develop differently, for instance perhaps altruism only emerges in societies that frequently win the inter-society competitions. There is no difference between the implementation of the societies so their labelling is arbitrary, there is no sense in presenting average results for “society A” as society A in the first simulation may have nothing in common with society A in the second simulation. Instead the societies are organised by how successful they were in the intersociety competitions, as this is a good indicator of the effect that the intersociety competition manipulation will have had on that society. The results are displayed such that the frequencies of behaviours in the most successful societies are displayed on the left.

Table 5.3: Baseline results for inter-society competition by wins.

	Society One (Most Wins)	Society Two	Society Three	Society Four (Least Wins)
Wander	0.97	0.99	0.83	0.75
Forage	57.81	58.48	60.29	58.92
Steal	40.83	40.04	38.46	39.92
Share	0.38	0.49	0.42	0.41

Table 5.4: Baseline results for inter-society competition by losses.

	Society One (Least Losses)	Society Two	Society Three	Society Four (Most Losses)
Wander	0.90	1.05	0.81	0.80
Forage	58.16	58.46	58.79	60.08
Steal	40.52	40.03	39.98	38.72
Share	0.42	0.46	0.42	0.40

This presentation of results is provided for easy comparison with the replication results (see table 5.2), the average results are comparable to the “No Reputation” condition, although there is some variation between the different societies. These results present only a crude measure of how success in intersociety competition relates to the behaviours that emerge within that society. A society that obtained the most wins in one simulation might have won 26% of the time, while another had won every competition, these situations would produce different behaviours. A more accurate measure of the impact of intersociety competition on emergent behaviours can be obtained by examining the correlations between each society's level of success (measured in the frequency of wins and losses) and its behaviour (measured by how frequently each of the four behaviours occurs). The results of this analysis for the

baseline simulation is presented in Table 5.5.

Table 5.5: Correlation matrix for inter-society competition without reputation

	Wins	Losses	Wander	Forage	Steal	Share
Wins	-					
Losses	-0.83**	-				
Wander	0.46**	-0.4**	-			
Forage	-0.39**	0.3*	0.06	-		
Steal	0.33*	-0.25	-0.18	-0.99**	-	
Share	0.05	-0.08	0.12	0.17	-0.22	-

* $p < 0.05$ ** $p < 0.01$

Some of these relationships are trivial, but confirm that the simulation is working correctly; for instance, the cause of the negative relationship between wins and losses should be obvious. The presence of societies that alternate between being highly successful and highly unsuccessful (e.g. Region four in Figure 5.2) leads to some societies having a high number of both wins and losses, but in the majority of cases societies that win frequently do not lose frequently and vice-versa. The negative relationship between forage and steal is also trivial. They both describe how often a behaviour is performed and the total number of behaviours performed is fixed (one per agent per cycle) so any significant increase in one behaviour will normally lead to a significant decrease in the other behaviours. No systematic patterns involving the wander and share behaviours can be observed. This is because rules that produce these behaviours are only occurring as result of the infrequent random mutations in inheriting internal rules, so these behaviours are rare and inconsistent across generations. In general correlations between winning and losing or between any of the four behaviours will

be trivial for the reasons described above, but they must still be examined as non-trivial relationships emerge between these variables in other conditions.

The important correlations in table 5.5 are those between the measures of success (wins/losses) and behaviours (wander/forage/steal/share) as these results show that inter-society competition has a systematic effect on the behaviours that emerge in the competing societies. The frequency with which societies gain or lose extra resources influences how frequently agents in those societies will wander, forage or steal. Competitive success leads to a decrease in foraging behaviour and an increase in wandering and stealing. When a society wins an inter-society competition, every agent within it gains some food. This makes foraging less necessary (as there is another source of food) and stealing a more attractive option (as other agents have more food to steal). It also explains the asymmetry between the effects of winning and losing competitions, losing intersociety competitions did not lead to significantly more stealing. After an intersociety success, the minimum amount of food gained in a successful steal action for any agent in that society goes up (to one plus the reward). However, after a loss the minimum amount of food gained in a successful steal action does not go down (as it cannot be lower than one, or the target agent would have died, making it an invalid target for the steal action).

Intersociety Competition: Friendly Forever

In the no-reputation condition introducing inter-society competition does not significantly alter the level of altruism in the simulation, as this condition provides no individual or group level advantage to altruistic agents. The next step is to examine how inter-society competition influences the condition in which Spronck & Berendsen (2009) found that altruism occurred (the "friendly forever" condition) to determine what effect inter-society competition has on the emergence of altruism (see Table 5.6). In this condition agents gained a positive reputation from a single act of sharing, regardless of their stealing behaviour.

Table 5.6: Correlation matrix for inter-society competition for "friendly forever"

	Wins	Losses	Wander	Forage	Steal	Share
Wins	-					
Losses	-0.65**	-				
Wander	0.01	-0.08	-			
Forage	0.01	-0.38*	-0.11	-		
Steal	0.02	0.32*	0.13	-0.98**	-	
Share	-0.08	0.43**	-0.02	-0.83**	0.66**	-

* $p < 0.05$ ** $p < 0.01$

Predictably there are negative correlations between winning and losing and between common behaviours in the simulation. However, there is a positive relationship between stealing and sharing. The positive relationship between these behaviours is probably stronger than the correlation coefficient of 0.66 suggests, as there is an inherently negative relationship

between any two behaviour outcomes; this is because the total number of behaviours exhibited is fixed, as described above. The high positive relationship between sharing and stealing behaviour can be explained as a result of the emergence of punishing agents. The way in which these agents use the stealing behaviour creates an environment in which more agents can successfully adopt the sharing behaviour, which leads to a positive correlation between sharing and stealing behaviours.

Under the friendly-forever condition the number of competitions that an agent's society wins does not influence that agents eventual behaviour, as was the case in the baseline simulation. Instead the number of losses is more influential. Agents in simulations that fare poorly in the inter-society competitions steal and share more frequently and forage less frequently. The friendly forever condition encourages the emergence of punishing agents and altruistic agents; these agents primarily use the stealing and sharing rules. The overall fitness of a society is determined by the total health of agents in that society, which is a reflection of how much food they have managed to consume in total. The steal and share actions do not affect the fitness of the society as they only move health around within that society, but the forage action increases the chances of gathering new food, improving the societies total health. So while the friendly forever condition encourages the emergence of altruism, the societies which develop it are less successful. This result cannot be used to explain the paradox of altruism as multilevel selection theory (e.g. Wilson & Sober 1994) would predict that at the group selection level, less successful groups would be selected against, this point will be picked up again in the general discussion.

Intersociety Competition: Basic Reputation

While the friendly forever condition produces a level of altruism, it is not a realistic representation of reputation. Few, if any, people would positively evaluate an individuals who consistently mistreated them on the basis of a single act of kindness. The basic reputation condition in Spronck & Berendsen (2009) is a more realistic implementation of reputation, as agents form an opinion of each other on the balance of their positive and negative behaviours. In the original study altruism did not emerge in this condition, but the introduction of inter-society competition causes some societies to develop higher levels of altruism (see table 5.7).

Table 5.7: Correlation matrix for inter-society competition for the basic reputation condition.

	Wins	Losses	Wander	Forage	Steal	Share
Wins	-					
Losses	-0.82**	-				
Wander	0.03	0.07	-			
Forage	-0.65**	0.55**	-0.4**	-		
Steal	0.66**	-0.56**	-0.37*	-0.99**	-	
Share	0.41**	-0.37*	-0.13	-0.47**	0.45**	-

* $p < 0.05$ ** $p < 0.01$

Many of these relationships are similar to those obtained under the "friendly forever" condition. Winning is still negatively correlated with losing, the behaviours are still negatively correlated with each other, except for stealing and sharing which increase together due to the co-evolution of punishing and altruistic agents. However under these parameters the

competition between societies has played a much greater role in shaping the agents behaviours than it did in the friendly forever condition. Other than wandering behaviour (which is minimal in all of the societies in this condition) the frequency of all of the behaviours is related to how frequently the societies win or lose the intersociety competitions.

Compared to the “friendly forever” condition, the role of intersociety competitions with respect to altruism have been reversed, under this condition increased altruism is now correlated with greater intersociety success. This is because the influence of punishing agents is fundamentally different under the “basic behaviour” condition. Under the “friendly forever” condition a punishing agent would not steal from any agent whose history contained at least one act of altruism. This made them more likely to steal from younger agents who had a shorter history. Conversely under the “basic reputation” setting the punishing agent would refrain from stealing from agents whose history did not yet contain more selfish than altruistic acts, so newborn agents were never targeted and older agents were more likely to be stolen from. This meant that in the “basic reputation” condition, the population that produced altruism also had more agents reach mating age, which in turn lead to a larger population. A larger population made it possible for the society to acquire food from more spaces at once, leading to more success in intersociety competitions, despite performing fewer 'forage' actions to actively move agents to areas containing food (the forage action was not necessary for an agent to consume food on its space, it simply caused the agent to move to a space containing food).

Under these conditions the level of altruism is significantly changed by how well the groups fare in the inter-society competition. This analysis is not quite sufficient to support the hypothesis that the introduction of multiple-identities allows the basic reputation condition to produce significant altruism. For this, the level of sharing observed in the "multiple societies

and basic reputation" condition is compared to the level of sharing observed in the "multiple societies and no reputation" condition by means of a t-test. Where multiple societies are present, the basic reputation setting produces increased altruism $t(86) = 4.22, p < 0.001$. While the absolute level of altruism observed in this condition is still low ($m = 0.75$ s.d. = 0.21), it is proportional to the level of ingroup competition that produced it, which is also comparatively low (1 unit of food per cycle, compared to 10 for foraging or 25 for stealing). This supports the claim that simulating additional social groups produces novel results, by changing the impact of an existing parameter.

Discussion

ICEM expanded upon existing research into the conditions in which altruism occurred, finding that it is possible for altruism to occur in a more realistic reputation condition than that presented in Spronck and Berendsen (2009). When intersociety competition is introduced altruism is able to occur under a reputation condition in which an individual's reputation is based upon a balanced account of all of their actions, rather than individuals obtaining a positive reputation through a single act of altruism. The analysis of ICEM is also detailed enough to determine the cause of altruistic behaviour emerging under this condition, highlighting the role of punishing behaviours in the emergence of altruism. While this behaviour likely occurred in the original Spronck and Berendsen (2009) friendly-forever condition, they did not identify it in their analysis. The analysis of ICEM showed that the punishing agents were responsible for the emergence of altruism. Punishing agents identified other agents with a low reputation and stole their resources, limiting their ability to mate and pass on the behaviours that lead to them having a low reputation. The presence of punishing

agents made agents who had obtained a high reputation through engaging in altruism more likely to survive and pass on their behaviours, leading to altruistic behaviours becoming more common in the simulation. This advances existing agent based modelling literature on the role of punishing behaviour, for example Ord's (2006) prisoners dilemma model. This study demonstrates the effectiveness of agents who punish defectors as well as agents who moderate their punishments based on whether the defector was themselves engaging in punishing behaviour. ICEM is not sophisticated enough to produce the latter behaviour, but enriches research upon punishing agents by demonstrating how they evolve in a population that does not contain many altruistic individuals.

The impact of the punishing agents varied depending on whether the observable reputation was based upon a single past action, or several and whether agents were considered friendly by default or not. These results predict that the type of information available for forming judgements would have a strong influence on the nature of social punishments in a population. Additionally, the punishing behaviour and altruism that emerged was only adaptive for the groups that it emerged in when reputations were formed using more complete information. In the friendly-forever condition (in which a positive reputation was obtained through a single altruistic act) societies containing more altruism fared poorly in the inter-society competitions. Multilevel selection theory (e.g. Wilson & Sober 1994) would predict that such societies would eventually fail and the behaviour would be lost. However the altruism that emerged in the more realistic basic reputation condition occurred in the most successful groups, demonstrating that the new altruism emerging in ICEM is consistent with existing theories on the emergence of altruism.

In the context of the evolution of altruism, the results of this simulation show a need for a more holistic approach to studying these factors. ICEM revealed that the relationship

between the reputation and multi-level group selection accounts are complex, as switching between two different types of reputation reversed the impact of success in intersociety conflict. It would be profitable to develop a complete model of the evolution of altruism that showed the interaction between the most commonly proposed factors.

Expanding Spronck & Berendsen (2009) to consider more than one society fundamentally changed the findings of the original experiment. The “basic reputation” condition was able to produce altruism in the presence of multiple competing societies, suggesting that a realistic interpretation of reputation can be critical to the emergence of altruism. Additionally their original finding that conditions that might promote altruism lead to a reduction in stealing behaviours can be questioned on the basis that ICEM demonstrates that a form of stealing behaviour (punishing) and altruism co-evolve. Together with the findings of the simulations presented in the previous chapter this produces a strong argument for modellers becoming more conscious about the implications of their decisions concerning the number of social groups modelled.

Chapter Six: Conclusions

General Discussion

In the previous chapters, I have presented five agent-based models developed to address questions from areas of social psychology including: Intergroup contact, group dynamics, segregation and altruism. While some of these domains are related, they are sufficiently different that generally, the implications of studies from one area do not strongly influence the implications of another. The commonality between the studies presented in this thesis is the agent-based modelling method; these models both demonstrate the usefulness of agent-based modelling to psychologists across multiple domains, and they can identify ways in which the agent based modelling method can be improved in future research. Therefore, this chapter reviews theoretical contributions of each of these models, and then to look at the methodological features of these ABMs, and concludes by drawing out the contributions to advancing our understanding in the field of group processes.

Theoretical Results

To begin, each of the models was developed to study issues in a different theoretical area, the destigmatisation simulation (DSIM) examined the contact hypothesis (Allport, 1954), the dynamic groups model (DGM) and its extension looked at fixed social group size (Hill and Dunbar 2002), the multiple group segregation model (MGSM) was concerned with

expanding a segregation model (Schelling, 1971) and the intergroup competitive evolution model (ICEM) sought to expand a model to explain the paradox of altruism (Spronck & Berendsen, 2009). In addition to the contributions each model made to its respective theoretical area, most of the models also had implications for each others areas. These contributions are summarised below in Table 6.1.

Contributions to each of these theoretical areas shall be assessed in turn. Each section will briefly reintroduce the problem and then describe the contributions of the main simulation addressing that area. Then it notes any implications that the other simulations have for that area before drawing the findings of that area together.

Table 6.1 Theoretical contributions of the models.

Theoretical Area	Simulation	Result
Contact Hypothesis	DSIM (chapter 2)	<p>Self fulfilling prophecies produce the contact effect.</p> <p>The contact effect saturates at higher levels of contact.</p> <p>Contact moderators impact overall improvement.</p> <p>Contact moderators do not always facilitate each other.</p>
Group Size	DGM (chapter 3)	<p>Group size dynamics can be understood in terms of conformity, ingroup rejection and rejection-identification.</p> <p>Social groups can conform to fixed group sizes, a property previously observed in personal social networks.</p> <p>Social group propagation requires the level of ingroup rejection to vary, either over time or between individuals (though see below)</p>
Group Size	E-DGM (chapter 4)	<p>With multiple identities, social group propagation with a fixed level of ingroup rejection becomes possible.</p> <p>Some group memberships may be incompatible with holding large numbers of social identities.</p>
Segregation	MGSM (chapter 4)	<p>Mild preferences for like neighbours still lead to higher segregation with multiple groups. However this effect diminishes as more groups are considered.</p> <p>Individuals who belong to multiple groups decrease segregation between those groups.</p> <p>When the above occurs it can increase the social instability of a third, unrelated, group.</p>
Segregation	DGM and E-DGM	<p>The level of segregation experienced by a group is influenced by ingroup rejection as well as by the preference for like neighbours.</p> <p>Low levels of ingroup rejection can produce high degrees of segregation, even where the preference for similar agents is low or non-existent.</p>
Altruism	ICEM (chapter 5)	<p>A combination of reputation and intergroup competition leads to the emergence of altruism.</p> <p>This occurs through the medium of “enforcement agents” that promote altruistic behaviour by punishing selfish agents.</p>

The Contact Hypothesis

The contact hypothesis predicts that contact between members of different social groups reduces intergroup bias, and diminished the negative effects of exclusion including the level of perceived stigma (Allport, 1954). Allport argues that there were four optimal conditions that either facilitated, or are needed to facilitate the reduction in bias. Since its publication, contact theory has been applied to a broad range of issues including regional differences in prejudice (Wagner, van Dick, Pettigrew, & Christ, 2003) and the educational main-streaming of physically and mentally disabled children (Naor & Migram, 1980). The effects of contact in reducing stigmatisation have been shown to apply in longitudinal studies (Christ et. al., 2010).

However, a number of questions remain, including whether the optimal conditions are really necessary. For example, Amir (1969) suggests that there might be two additional conditions for contact (intimate contact and pleasant contact). As a result, over the past few decades researchers have continued to explore the role of several contact moderators. This research has been systematically reviewed in a meta-analysis by Pettigrew and Tropp (2006) (also see introduction of chapter 2). A shortcoming with the studies available for this meta-analysis, however, is that they only considered Allport's optimal conditions as 'all present' or 'all absent', rather than investigating them individually. This means that it is not possible to tell what the relative weighting of the individual moderators actually is. The DSIM, chapter two, was able to produce a 'contact effect', that is a reduction in intergroup bias as a result of increased contact between groups.

While the intergroup literature has mounting evidence to support the notion that increased contact increases acceptance and decreases stigma, there is relatively little evidence

to suggest that there is a link between self fulfilling prophecy and this contact effect. However, a contact driven reduction in stigma is one of the properties that emerges from DSIM's implementation of self-fulfilling prophecy theory (e.g. Jussim, et al. 2000). This allows DSIM to demonstrate a theoretical link between SFP and intergroup contact. In conditions where the other assumptions of the model hold true, this allows SFP to act as an explanation for why the contact effect occurs.

In addition to expanding on how this effect might occur, DSIM also reveals features of this effect that may have been overlooked in previous research. The destigmatisation effect modelled by DSIM saturated. In other words, over time the groups' opinions of other social groups converges so destigmatisation would cease irrespective of intervention. But, the result indicate that increasing intergroup contact through interventions would only be effective for a limited time, after which they would lead to no further decrease in stigmatisation.

This finding is also important for the contact moderators, as previous research has only identified which moderators lead to an overall reduction in stigmatisation (e.g. Christ et. al. 2010). It has not been able to distinguish between those moderators that cause a greater overall reduction in stigmatisation, and those that merely cause the destigmatisation to happen more quickly. DSIM is able to distinguish between these and found that only moderators that were implemented as influencing the quality of the interactions, rather than their frequency or the magnitude of their effects, produced a greater improvement in the long term. This set of findings maps on well to Allport's (1954) finding (reproduced in Schwartz & Simmons, 2001) that the quality of the intergroup interactions was more important than their quantity for reducing stigma.

DSIM's final contribution to this area of the literature is to start to unravel the complicated relationship between several important contact moderators; that is, to explore the

conditions under which they reinforce each other, as well as those in which they impede each other at reducing bias. A great deal of previous 'contact research' has only compared a condition in which all moderators are present with one in which all moderators are absent as discussed. The systematic interaction between moderators is rarely studied. DSIM fills this gap testing the moderators independently and together; the findings illustrate that 'optimal conditions' all have a facilitatory effect upon reducing bias through increased ingroup contact. It also highlights that they impede each others positive effects upon how quickly bias is reduced, but enhance each others impact upon how much bias reduction occurs. Though moderators that only influence the quantity of the interactions rather than the quality of them that they will influence only the speed of destigmatisation for the stigmatised group. While their facilitatory effects make it easier to find effects for contact moderators by combining them (Pettigrew & Tropp, 2006), it is important to consider them individually as they have different long term effects – with best effects for interventions being driven by the quality not the quantity of contact.

The other models presented in chapters 3-5 do not contribute to this theoretical area as they contain no explicit measure of stigmatisation and the contact hypothesis is concerned with changes in stigmatisation. The degree of segregation found in several of the models can be viewed as an implicit measure of stigmatisation because reduced social links and denial of access to social resources are one measure of stigmatisation. However the degree of segregation would also be the only measure of how much contact takes place within these simulations, thus any conclusions drawn would be trivial.

Considering DSIM's findings alongside other models starts to build a more complete view of stigmatisation. Several of the models, most notably the dynamic group model (DGM) and its extension, examine the conditions which cause social groups have a greater or lesser

amount of contact with each other. Taken together these models form an account of the conditions under which various levels of contact might be obtained and the likely effects of that upon intergroup stigmatisation. These implications will be discussed under the domains relevant to the appropriate models.

The main finding of this thesis in terms of the contact hypothesis is that Allport's (1954) optimal conditions all have a facilitatory effect upon destigmatisation through the contact effect, that they enhance each other and that if they only influence the quantity of the interactions rather than the quality of them that they will influence only the speed of destigmatisation rather than its ultimate level. While their facilitatory effects make it easier to find effects for contact moderators by lumping them together (As in Pettigrew & Tropp, 2006) it is important to consider them individually as they have different long term effects.

Social Group Size

While sociologists have speculated that social networks might have a fixed upper limit for some time (e.g. Pool and Kochen, 1978), their focus has been on estimating the size rather than understanding why such a limit might exist. In contrast, Hill and Dunbar (2002) proposed that the reason that a maximum social network size is obtained might be linked the neocortex size. However, other studies have suggested that this phenomenon might have a different basis. For instance Roberts, Wilson, Fedurek and Dunbar (2007) demonstrates that personality factors influence an individual's social network size. Another important question in this field is whether a similar size restriction might influence social groups, a concept highlighted by Dunbar (2008). Dunbar mentions the commonality of groups of 150

individuals and that is common in non-academic circles (e.g. Allen 2004). These questions are related, as some of the proposed explanations of social network size would not apply to a social group. For instance the biological account in Hill and Dunbar (2002) could not apply to a social group as such an entity does not have a neocortex size. As such when a fixed social group size is observed, a new explanation is required for why it should occur.

The dynamic groups model (DGM) presented in Chapter 3 addresses some of the questions surrounding this debate, by successfully modelling a social group with a fixed size. The DGM is based on a social psychological perspective that identifies a few key variables that influence an individual's prospect of leaving and joining social groups. In Chapter 3 I review social identification as a framework facilitating motive for group membership. Social rejection (Jetten, Branscombe, Spears & McKimme, 2003), conformity (Nail & Helton, 1999) and the rejection-identification hypothesis (Branscombe, Schmitt & Harvey, 1999) were all critical in DGMs ultimate success in modelling a social group with a fixed size.

Under all conditions in the DGM (Chapter 3) and most conditions in the extended DGM, social groups obtained a fixed group size. As well as demonstrating that a fixed group size for social groups is possible, the DGM findings would predict that fixed group sizes are obtained due to the behaviours that were implemented in the model. That is to say that social conformity and ingroup rejection act in opposition to each other and as long as the processes remain fairly constant (i.e. the level of ingroup rejection did not change and the number of individuals required to cause conformity did not change) the group would also obtain a fixed size.

DSIM (Chapter 2) and MGSM (Chapter 4) are both models that deal in fixed group sizes and have little to contribute to a discussion on group sizes. However ICEM (Chapter 5) was based around examining four competing groups that changed in size. In this model

groups sizes were primarily determined by the availability of physical resources and the effectiveness of the agents behaviours in terms of utilising them; groups that acquired more resources by winning intergroup competitions or more effectively harvesting those appearing randomly in their environment became much larger. However a fixed group size did not occur, as a population that was too successful could consume all of its food and start to starve, causing the group to experience a series of population surges and crashes of increasing magnitude, until a crash brought the population low enough to no longer be a consistently successful group. Groups with rapidly changing sizes in this model were unstable and often suffered for it, so this can provide an evolutionary explanation of why behaviours that lead to more stable group sizes might be observed in human populations in the first place.

The main finding of this thesis in terms of social group size is that social groups can obtain a fixed group size, in a manner similar to social network sizes predicted in Hill and Dunbar (2002). This static group size can be explained as a tension between factors that drive a group apart (such as ingroup rejection, Tajfel & Turner 1979) and factors that unite a groups members (such as conformity, Bond 2005) with a static group size only being obtained where the influence of these factors on the group remains relatively constant.

Segregation

Segregation describes a social issue whereby social groups in society are divided, with few interactions between the groups and typically with unequal access to social resources such as education, housing and jobs (e.g. Pettigrew, 2008). Schelling (1971) modelled racial segregation between black and white Americans, looking at why some cities self organised to

have 'black' and 'white' areas. The conclusions of the simulation can be generalised to other examples of segregation. This model is widely regarded as a critical model in the development of agent based modelling as a field. Schelling investigated why there were high levels of racial segregation occurring in locations where nobody had a particularly strong preference for racial segregation. The model found that when people had mild preferences for individuals who were similar to them, it could lead to high levels of segregation. The multiple group segregation model (MGSM) presented in Chapter 4 is an extension of Schelling's work and it develops the agent based modelling methodology by looking at the impact of simulating multiple social groups, rather than simply modelling a two groups. This extension is driven by findings from social psychology concerning the ways in which people can hold multiple group identifications and/or affiliations at any given time (e.g. Crisp, Ensari, Hewstone & Miller, 2002).

In MGSM the key finding that segregation can emerge from small preferences for similar neighbours becomes less pronounced. A small level of ingroup favouritism still produced segregation, but the inclusion of the third group resulted in a lower level of segregation. Examining the cause of this makes it clear that this is an effect that increases as more groups are involved, as the greater number of alternative social groups increases the opportunities that the agents have for outgroup contact. As individuals often have a great many identities this might seem to imply that the relationship Schelling (1971) uncovered is unable to explain the degree of segregation that is observed through minor preferences. However if the agents had additional identities that did not influence their decisions to move this effect would disappear again, so in fact the model suggests a notion that segregation can result from minor preferences for segregation only so long as few identities are salient.

Additionally, when MGSM is used to model a situation where only some groups had

access to agents with multiple identities produced novel results. As would be predicted from the literature on the cross-cutting effect (e.g. Fry, 2009) the introduction of multiple identities decreased segregation (one form of stigmatisation) between the groups that had these cross-cutting agents. However it also served to increase segregation for members of a third, unrelated, group. Thus MGSM predicts that interventions around nurturing better relations between members of two conflicted groups may have negative consequences for members of other social groups, a possibility that has implications for how such interventions are handled.

The DGM (Chapter 3) and E-DGM (Chapter 4) make direct contributions to the segregation domain. The degree of isolation experienced by the target group was measured in all simulations and examining this data suggests some interesting interactions between the level of positive ingroup interactions within a group and the level of segregation that the group experiences. In particular that a relatively low preference for similar neighbours is not sufficient to produce a low level of segregation, as might be predicted solely from the segregation model presented above. Instead a low level of segregation was only produced where a low preference for similar neighbours was coupled with a relatively high level of positive ingroup interactions.

Changing focus, the models presented in Chapters 2 and 5 do not address how segregation occurs as in chapters 3-4, but rather demonstrate the consequences of segregation. DSIM (Chapter 2) demonstrates that as a community becomes more highly segregated, a higher level of stigmatisation is able to persist. This is consistent with the conceptualisation of social exclusion as a form of stigmatisation (e.g. Major & Eccleston, 2005), as the models demonstrate that social segregation implies stigmatisation. It also demonstrates that intergroup contact has positive effects for at least one of the groups involved, which provides part of the motivation for studying segregation. In chapter 5, ICEM demonstrates that segregation is an

important component to the emergence of altruism and the altruism that emerges in ICEM is only expressed towards other agents that belong to the same group, as behaviours emerged in which agents gave food to other members of their group, but no intergroup behaviours emerged that caused a group to sabotage its chance of winning the intersociety competitions to benefit another group.

As a whole the thesis shows that segregation is more likely to occur where only a limited number of identities are salient (or available) to relevant decisions, and where the dominant ingroup has a low level of ingroup rejection.

Altruism

The paradox of altruism describes one pattern of behaviour that has been a consistent problem for evolutionary theorists. The challenge is to explain why, if all behaviour is ultimately driven by mechanisms developed to replicate themselves, do individuals have mechanisms that drive them to altruistic behaviours. Over the years many researchers have investigated this problem and proposed a variety of solutions (e.g.s Hamilton, 1964, Trivers, 1971). The model presented in the fifth chapter extends an evolutionary simulation (Spronck & Berendsen, 2009) to further examine this problem.

The intersociety competitive evolution model (ICEM) produced a pattern of evolution from which a significant level of altruistic behaviour emerged. This occurred under the "moderate reputation" condition, in which an agent's reputation was determined by comparing how often they performed greedy and altruistic behaviours. By contrast, the original Spronck and Berendsen (2009) study which produced altruism only during the "friendly forever"

reputation condition, in which an agent's reputation was positive if it had ever performed a single act of altruism. The ICEM result is much more realistic, as it is trivial to observe that individuals assessments of each other do not ignore negative behaviours. Introducing only the moderate reputation condition or the intergroup competition condition alone did not result in any increase in altruistic behaviour, thus it is a combination of these factors that produces altruism in ICEM. ICEM's principle contribution to evolutionary psychology is to demonstrate that while reputation or multilevel selection may be individually unable to explain the emergence of altruism, in conjunction they can lead to evolutionary processes producing altruism. This line of argument can be extended to other proposed solutions to the paradox of altruism though further research is needed to examine these in conjunction. ICEM's main contribution to social psychology is to link the emergence of altruism to normative behaviour, as the emergence of altruism is facilitated in ICEM through the emergence of 'enforcement agents' that punish other agents for failing to take altruistic actions.

ICEM also demonstrates one of the reasons that normative behaviour has arisen in humans, through examining the mechanism that makes it possible for altruism to emerge. In the simulations in which altruism emerged, other agents developed a behaviour causing them to steal from greedy agents instead of altruistic ones, making altruism a more attractive choice for other agents and eventually increasing their own resources. This promotes the survival of altruistic agents and makes that behaviour more prevalent in the population. Taken with other lines of research (e.g. Ord 2006) this contributes to an argument that the most effective societies must be maintained by some number of agents enforcing group norms.

The other simulations, presented in Chapters 2-4 do not directly examine altruism. However there is a related concept that comes out of this line of research as a whole, that

systems as a whole are highly sensitive to relatively subtle effects on an individual level. This will be discussed in more detail in the methodological conclusions, but its relevance here is that the gap between the level of altruistic behaviour expressed in ICEM and the level of altruistic behaviour present in the world might be explained through an aggregation of the smaller effects observable in ICEM to a larger scale.

The principle contribution of this chapter is to demonstrate that altruism can emerge through a combination of factors, where one factor alone would not be enough to explain it. The findings of ICEM suggest looking for solutions to the problem of altruism by examining how the existing explanations interact with each other.

Methodological Results

The following section explores the features of the agent-based models as related to the models contained in this thesis. In some cases, the importance of these features coincides with existing research on agent based modelling, such as Smith and Conrey's (2007); however, in other cases methodological aspects of the studies offer an opportunity to improve the application of the agent based modelling method for future research. This review is organised about four major themes: theory building, sensitivity, steady state convergence and multiple social groups (see table 6.2).

Table 6.2 Methodological contributions of the models.

Feature	Model	Result
Theory Building	DSIM (chapter 2)	Self fulfilling prophecy can lead to destigmatisation via contact effect.
	DGM and E-DGM (chapters 3 and 4)	Identification, conformity and rejection-identification imply fixed group sizes in line with some (incorrect but interesting) interpretations of Dunbar.
	MGSM (chapter 4)	Multiple social group research implies that the cross-cutting effect (a) occurs and (b) isn't universally beneficial.
	ICEM (chapter 5)	The paradox of altruism can be explained by multilevel selection and reputation in conjunction.
Sensitivity	DSIM	Small degree of contact changes lead to big final stigmatisation changes. Small moderator changes can lead to big final stigmatisation changes (most notably status)
	DGM and E- DGM	Small changes in ingroup rejection lead to large group size changes, especially note the critical area at ~50% in DGM and ~60% in E-DGM
	MGSM	Mild preferences lead to high segregation (as per original) Mild presence of multiID agents between two groups can leave a third significantly more segregated
	ICEM	Very small intergroup competition change influences most results (most importantly altruism)
Steady State Convergence	DSIM	Opinions converge to a point, this process describes all of the destigmatisation results in the simulation.
	DGM and E- DGM	Group sizes converge to a point (this is a key finding, finding a common property between group size to social network size)
	MGSM	Looking at segregation over time, it converges to a point defined by the parameters of the model.
	ICEM	The property of convergence is a critical premise of the model.
Multiple Groups	E-DGM (chapter 4)	Including multiple groups allows for unstable social networks finding.
	MGSM	Including multiple social groups weakens original finding (segregation is less marked) Including multiple social groups allows for secondary finding of group ostracism via greater integration of two other groups.
	ICEM	Including multiple social groups allows for finding that reputation can work under reasonable conditions.

Theory Building

A theoretical agent-based model should be able to relate existing theories to each other, in order to generate a more complete picture of the relevant domains. As proposed in Smith and Conrey (2007) and detailed in the introduction the goal of theoretical agent based modelling is to develop a model from one set of findings and use it to explain a different set of observations, thus linking those two areas. The original models presented in Chapters 1-3 were each based on one or more psychological theories and had results in line with different theories, not included in their development. This allowed them to demonstrate links between previously unrelated concepts, furthering the understanding of both areas.

In DSIM an implementation of self fulfilling prophecies drove the emergence of the contact effect within that simulation. While the contact literature was used later in the models development to study various contact moderators, this did not take place until after the first study in which a contact effect emerged. The result contrasts with the existing relationship between stigmatisation and self-fulfilling prophecies as SFPs are generally used as an explanation of why stigma is maintained (Jussim et al. 2000). DSIM shows the theory in a new light, demonstrating one mechanism by which it might reduce stigmatisation in a way consistent with existing research into the contact effect.

The implementation of the DGM was based on social identity theory, conformity and rejection-identification. While some of these theories are in the same area as social group size, no previous study had drawn these theories together in order to provide a framework for the investigation of this topic. The DGM and its extension are a first step towards building a coherent framework, relating these theories to provide clear explanations for observed phenomena.

MGSM uses multiple identification research to extend segregation research to handle a situation in which individuals might be a member of multiple groups. This produced two results, the first was that segregation was weakened in the presence of agents that held multiple identities. The second is that a third group might experience increased segregation due to greater cooperation between two principle groups. This allows a conclusion to be drawn from the multiple identity literature regarding the possible consequences of cross-cutting groups. The conclusion itself is striking as previous cross-cutting research has focused on the benefits of individuals with cross-cutting group membership, without exploring possible side effects upon other groups.

ICEM served to relate the concepts of reputation and multilevel selection to each other in order to provide an explanation for the emergence of altruism. While previous researchers has typically looked at these concepts individually, ICEM examines two factors in conjunction. This produces a result that could not otherwise have been obtained, as neither of them are individually able to produce emergent altruism. Further research into the paradox of altruism must take a wider view encompassing several explanations that may not be sufficient individually.

The theory building properties of agent based models are often the reason that this approach is taken to study a particular domain. The ABMs in this thesis all make use of this property to contribute to their relevant domains, advancing psychological theory in several areas.

Sensitivity

Sensitivity describes the attribute that some models have where a small change in a behaviour conducted at a microscopic level (e.g. an individual persons behaviour) produces a large change at a macroscopic level (e.g. at the level of a society or culture), again the importance of the microscopic-macroscopic interactions are highlighted in Smith and Conrey (2007). It is important to understand when and why this occurs, as it allows models to take the sometimes modest effects found in traditional research and use them to explain much larger observable phenomena. Each of the models presented in this thesis demonstrated this principle at work on at least one occasion.

In DSIM, small changes in the amount of contact agents had lead to large changes in how stigmatised the minority group was at the end of each simulation. At low levels of contact the changes were more dramatic, with the largest difference being between groups that had minimal contact and those that had no contact at all. Additionally many of Allports (1954) contact moderators were implemented as having relatively small effects, such as weighting the opinions of the agents in a 3:2 ratio rather than a 1:1 ratio. Modest changes to these moderators, most noticeably status, produced large changes in the final level of stigmatisation.

The DGM and extended DGM found that there were situations in which a small change to the level of ingroup rejection produced large changes in most of the outcome variables, including group size, identification, segregation and network propagation. The model was particularly sensitive to small changes in ingroup rejection around a “sensitive region” at the level of ingroup rejection made a dyad (two agents interacting in isolation) stable. This was due to the fact that whether a dyad was stable determined the type of clusters

that could form, which in turn had a large effect on identification and group size.

MGSM is a replication and extension of a model that famously found that a very mild preference lead to extreme segregation (Schelling, 1971). This finding is replicated, even with multiple groups segregation emerged from a low preference for similar neighbours. MGSM also demonstrates another type of sensitivity, which occurred where agents that belonged to a particular two of the three groups were added to the three group simulation. This caused the group that the agents were not members of to become ostracised, despite the fact that the group itself was entirely unchanged from previous simulations. Revealing this type of relationship is a unique contribution that agent based modelling can offer to psychologists.

Finally ICEM makes a very small change to an existing model (Spronck & Berendsen, 2009) which produces very different results to the original. The inter-society competition effect that is introduced can provide an agent with a maximum of one unit of food per cycle, compared to ten for foraging or twenty-five for stealing. However the presence of the competition leads to important changes in the models behaviour, causing altruistic behaviour to emerge under the 'basic reputation' condition. As this condition is more realistic than the other condition under which altruism emerged, this small change allows the model to provide a much better answer to the paradox of altruism.

In each of these models a small change to one of the models parameters produced a large difference in the models outcome. Especially in Chapters 4 and 5, where the key finding of the models presented emerged from such a change. This highlights the effectiveness of the agent-based modelling approach in uncovering the non-linear relationships between variables as a means to explaining observed phenomenon.

Stable State Convergence

All of the models described in the previous chapters converged to a stable state, at least in terms of the relevant dependent measures. Most agent based modelling efforts involve finding a stable state (even if this state is an repeating pattern, such as an oscillation about a point). The ABMs presented in this thesis follow the pattern, each model described a process by which a simulation approached some final stable state. Studying the process of this convergence was important to understanding the theoretical implications of the models.

This is most notable in DSIM, where the time course of the model was explicitly examined in order to determine the effects of Allports (1954) 'optimal conditions' upon the contact effect. Examining opinion convergence identified how the optimal conditions influence the speed at which the model converges to its final state, which in this instance is the speed of destigmatisation. Explicitly examining the time course of the model provided the capacity to measure how quickly an 'optimal condition' produces a reduction in stigmatisation, as distinct from the magnitude of this destigmatisation effect. By de-confounding these two things DSIM was able to make additional theoretical contributions, notably the prediction that some 'optimal conditions' improve only the speed of destigmatisation, not its magnitude.

In the DGM and the E-DGM the time course was not explicitly studied. However interpreting the results required an understanding of how the agents behaved over time. Chapters 3 and 4 include snapshots of how each model behaves at regular intervals, to help explain how the agents' behaviour changes over time, before reaching a final state. This made it possible to develop explanations such as why the results were so sensitive to changes in positive ingroup interactions around the 50-60 mark. Additionally, that the model converges to a particular result at all is an important finding in its own right, as it allows the model to

confirm its main hypothesis: That something similar to Dunbar's number (Hill and Dunbar, 2002) applies to social groups.

In the original Schelling (1971) model examining the time course of the model is not particularly important, it does not change the interpretation of the results. Though perhaps it serves to reassure the reader of the legitimacy of the model's counter-intuitive result. In MGSM it serves this purpose, but also helps to develop an explanation of a finding unique to this extension. Specifically that the presence of two integrated groups have an impact upon a third unrelated group. In this model agents from the third group often cannot find a place that they are satisfied with the number of like neighbours they have, as the other two groups have broadly integrated, leaving insufficient space for the third group to have an area of its own.

The principle that a model will converge towards some systematic behaviour from a random starting state is the cornerstone of evolutionary computing. In all successful evolutionary computing, for commercial and technological purposes as well as in research models, agents converge towards some optimal point over generations. It is not possible to show the importance of convergence to ICEM specifically, as it is so fundamental to the nature of the model that it is difficult to see how the research could have been conducted at all without acknowledging this feature of modelling.

It is important to remember that it is not inevitable that the results of a simulation will converge to some stable state. This occurred several times in the models presented in previous chapters. For instance, under some parameter settings the extended DGM did not reach a stable state. Also, a stable state isn't necessarily one static value, it might be that some equilibrium is reached and while change still occurs in the model the overall result is no longer changing over time. In some of the models, particularly the DGM and the evolutionary model, this process of convergence was critical to the models success, however in all of them

an examination of how this process occurred over time resulted in a much clearer understanding of the meaning behind those models results. In agent-based modelling research exploring how the model approaches its final state is as important as exploring what the final state might be.

Multiple Groups

Gotts and Polhill (2010) demonstrated that unstated assumptions in modelling can lead to misleading results. In their study modelling different areas (a variable previously thought to have had no effect) qualitatively changed the outcome of a simulation. Agent-based models frequently use the minimum number of social groups necessary to test their hypothesis, this aids in the theory building dimension discussed above. The fewer variables a simulation alters, the easier it is to draw conclusions about what caused a particular behaviour to emerge. The importance of modelling the minimum number of variables is discussed at length in Smith and Conrey (2007). The final three models, presented in Chapters 4 and 5 were aimed at considering the implications of a different assumption. In this instance the research aimed to see if always using the minimum number of social groups was influencing simulation results. Three existing models were extended to include additional social groups, drawing from the multiple group literature. These were Schellings (1971) classic model of segregation, the DGM presented in Chapter 3 and an evolutionary simulation by Spronck and Berendsen (2009).

The Schelling extension, MGSM, demonstrates that the core finding of the original Schelling model is reduced in the face of multiple social groups. Studying the reasons that this

occurs in the model it is reasonable to conclude that it would disappear entirely in the face of a realistic number of identities. As well as modifying the core finding this extension also allowed a number of novel observations to be made regarding the impact of multiple identities on segregation. For example the finding that while agents that belonged to multiple groups reduced segregation between the groups they belonged to, they ran the risk of increasing segregation for unrelated groups. For this model considering additional social groups altered the existing findings of the model and allowed it to explain additional phenomena.

Expanding the DGM to account for multiple groups resulted in many of the simulations reaching a different equilibrium. In these simulations the structure of the social network represented in the simulation was unstable and agents regularly changed who they interacted with. The core finding was maintained however, as in all conditions still reached a fixed size. The extension suggests a caveat to these findings, as in some situations the group can encompass all of the agents in the simulation. This does not represent a stable group size, as if the simulation had more agents the group would continue to grow, so it is a limitation of the simulation rather than an implication of a theory that causes the size to stabilise. In this case expanding the model in this way revealed a limitation upon the original model's findings as well as allowing it to explore a different structure of social group.

Examining the impact of multiple competing groups in evolution has long been suggested as a reason for the emergence of altruism and other traits (Wilson & Sober, 1994). ICEM examines this by explicitly modelling multiple groups and indeed the presence of multiple groups produces a higher level of altruism. However this only occurs when agents are able to form reputations and interact with each other based upon these reputations, neither property can individually explain the emergence of altruism. Extending this model to consider additional social groups alters the findings in a way that enables the model to better achieve

its original goal, finding an explanation for the paradox of altruism.

Overall this thesis provides clear evidence that expanding simulations to examine multiple social groups often changes the original findings of the model concerned as well as enabling new lines of investigation. This is observed in a classic “toy model” (MGSM), in a more complicated model (DGM) and a different class of model (ICEM). These results strongly suggest that future agent based modelling research should seriously consider the impact of this often overlooked parameter.

Limitations and Outlook

The agent based modelling paradigm has a lot to offer social psychology, as evidenced by the specific contributions of the models presented in this thesis (described above) and the general advantages of the approach (discussed in the introduction). The method has some notable limitations, such as generalisability, oversimplification and the generation of needlessly convoluted theories. Most of these can be overcome through complimentary research or appropriate simulation design.

A particular ABM might produce results which cannot be generalised to the population as the result is specific to the assumptions made in that ABM. This occurs most frequently where the results from an ABM are an artefact of a specific implementation of an assumption, rather than being results driven by the assumption itself. This occurred in chapter two, where the overall stigmatisation reduction result of the authority moderator was driven by the fact that authority agents were taking up space on the grid, rather than by intergroup interactions being sanctioned by relevant authorities. This problem can be addressed, as it was in Chapter

two, by looking to the explanations for how a particular result is derived in order to identify and reject such findings. It can also be overcome by modellers replicating each others work with slightly different implementations to ensure that their findings are robust (as in chapters 4-5).

Another potential pitfall is that a result might not be applicable to human populations due to the simplifications made in order to implement an ABM. It is necessary to simplify to generate an effective ABM, failing to do so creates other problems such as generating ABMs that rely upon so many theories that it is impossible to use the findings of the ABM to make a clear theoretical statement (see below). The decision of which facets of the domain being simulated to leave out is critical to the ultimate effectiveness of ABM research and should be guided by the quality of existing research available to use for a particular set of assumptions. Ultimately there is no way to guarantee that the decisions made in the design of a particular model are effective so this problem can only be overcome by using traditional research methods to confirm the findings of ABM studies. This occurs in chapter one, which contains a survey confirming the observation that a groups size and level of stigma experienced are related, but all of the chapters suggest further research that could confirm and elaborate on their findings. Note that ABMs still play an important role as part of larger research programs as they allow researchers to identify theoretical links that would otherwise be missed and thus can guide traditional research to more profitable areas.

Finally there is a risk that theory building ABMs that require multiple theories to function can produce theoretical findings that are not useful due to their complexity. A model that has a multiple assumptions driven by many different theoretical constructs will end up drawing conclusions that state that each of these assumptions is necessary to produce a specific effect. However, this is unlikely to be the case, in an overcomplicated model a

number of the assumptions are likely to be irrelevant to a particular outcome.. In any case the multiple requirements of such conclusions may wind up being too specific to be useful and very difficult to confirm through traditional research methods. On a practical level over-simulation can also produce a multi-factorial design that requires a very long time to simulate all possible combinations. This is avoided through adopting a minimal modelling approach (i.e. modelling the minimum number of variables necessary to produce a particular behaviour). However this can lead to its own problems, as demonstrated in chapters four and five, if relevant parameters are implemented for practical purposes but never tested. A frequently used, but often unreported compromise is to perform pilot studies that vary such parameters to ensure that they do not influence the main findings of the simulation.

This thesis demonstrates that the agent based approach is able to contribute to several fields and as technology improves it is going to become an increasingly useful approach. The improvements made to existing models in chapters four and five show the usefulness of extending existing agent based models using findings from multiple social identity studies. This implies that in addition to modelling more domains, it is profitable for future researchers to extend existing models accounting for this feature of social psychology as well as to look for other well established psychological theories that have been overlooked in modelling to look for further useful extensions to existing modelling research.

Final conclusions

The simulations presented in this thesis successfully use the agent based modelling paradigm to make theoretical contributions to social psychology: Self fulfilling prophecies are

sufficient to produce a contact effect; the contact effect saturates at higher levels of contact; contact moderators can influence the speed of destigmatisation rather than the final level; contact moderators do not always facilitate each other; group size dynamics can be understood in terms of conformity, ingroup rejection and rejection-identification; social networks can conform to fixed group sizes in a similar fashion to personal social networks; social group propagation is dependant upon a changing group dynamic or the presence of multiple social identities; membership with groups that have high ingroup rejection is incompatible with holding large numbers of social identities; the effect that mild preferences for like neighbours leads to segregation is weakened in the presence of multiple identities; individuals with multiple social identities play a role in reducing segregation between the groups they belong to, but may increase the isolation of other unrelated groups; segregation can be produced by low levels of ingroup rejection rather than a level of preference for similar neighbours; combining the reputation and multilevel selection accounts of the evolution of altruism produces altruism under conditions in which neither of the accounts could individually explain its emergence; altruism occurs through the emergence of enforcement agents, tying it in to explanations of normative behaviour.

In producing these theoretical findings the simulations have not only confirmed useful features of the ABM paradigm that had already been identified, but also show that in the specific arena of social psychology the paradigm can be usefully adapted through a focus upon the impact of multiple social identities. This highlights the growing potential of agent based modelling as a field in psychology and starts to drive us towards finding ways to adapt a method that has been so successful in other fields to be more useful to psychologists.

Appendices A: Survey Wording

Please rank each of the following groups based on how socially stigmatised they are within UK society. Each number (from 1-23) can only be used once. A higher number indicates that a group is more stigmatised, so if you believe that "Arabic people" is the most stigmatised group on this list you would record a 23 in the space next to them. Note: Social stigma is a severe disapproval of a group based on personal characteristics, or beliefs that are perceived to be against cultural norms of the mainstream society. Keep this in mind while you are making your selection.

Arabic people
Bangladeshi people
Black caribbean people
Black african people
Chinese people
People with cancer
People who are registered as disabled
Epileptic people
People over 60 years old.
Gay people
Homeless people
Immigrants
Indian people
Pakistani people
People serving a jail sentence
Jews
Muslims
People of mixed race
Poor people
Refugees
People registered as unemployed
Victims of crime
Women

Thank you for taking part in this study.

Appendices B: Source Code

This appendices contains the source code for the simulations described in chapters two through five. The simulations were originally written in netlogo, in version 4.0.5 and 4.1.3, but have been automatically updated to version 5.0.3. Netlogo user interfaces are not created in the code tab, if a variable is assigned before it is declared the reader can safely assume that it was manually set on the user interface before a simulation was started. The simulations were run by a call to a setup function, followed by a 'forever' call to a main function, which would repeat until a “stop” command occurred. In most cases the simulations output a text file that can be read by MATLAB for the purposes of statistical analysis rather than performing the statistics directly. This is intended as a reference only, there should be sufficient information to replicate each simulation in the appropriate chapter.

DSIM

```
;Group 1 is the majority, group 2 is the minority

globals [poscount negcount cycle dimension prevvalues prevplot]
turtles-own [g1opin g2opin group interest sanctioned-interactions]

;;Replaces setup simulation with a file
to file-setup
  file-open "config.txt"
  let settings file-read
  set agent-numbers item 0 settings
  set agents-in-minority item 1 settings
  set interaction-range item 2 settings
  set authority-figures-on item 3 settings
  set authority-figures item 4 settings
  set status-effect item 5 settings
  set status-effect-magnitude item 6 settings
  set co-operation item 7 settings
  set co-op-freq item 8 settings
  set common-interest item 9 settings
  set common-interest-perception item 10 settings
  set cyclestorun item 11 settings
  file-close
  setup-simulation
end

to setup-simulation
  ;;Clear previous simulation away
  ;; (for this model to work with NetLogo's new plotting features,
  ;; __clear-all-and-reset-ticks should be replaced with clear-all at
  ;; the beginning of your setup procedure and reset-ticks at the end
  ;; of the procedure.)
  __clear-all-and-reset-ticks
  set prevvalues []
  set prevplot 0
  set cycle 0
```

```

set poscount 0
set negcount 0
;;Create world sized for agents and groups
set dimension ceiling(sqrt(agent-numbers))
let height 0
let width 0
let placedagents 0 ;Note this is total agents placed so far
let outgroupagents agents-in-minority ;Note this is outgroup agents left to place
;;I think I've confused the variable names for width and height in this algorithm
;;It doesn't matter (since it's a square and therefore symmetrical on a 45 degree axis :P)
;;However I will need to be careful if this section needs changing later
while [width < dimension]
[
  while [height < dimension]
  [
    ask patch height width
    [
      set pcolor white
      if placedagents < agent-numbers
      [
        ifelse outgroupagents + placedagents = agent-numbers
        [
          sprout 1 [set color blue set shape "person" set g1opin 0.9 set g2opin -0.9 set group 2]
          set outgroupagents outgroupagents - 1
        ]
        [
          ifelse outgroupagents = 0
          [
            sprout 1 [set color black set shape "person" set g1opin 0.9 set g2opin -0.9 set group 1]
          ]
          [
            ifelse random(agent-numbers - placedagents) <= outgroupagents
            [
              sprout 1 [set color blue set shape "person" set g1opin 0.9 set g2opin -0.9 set group 2]
              set outgroupagents outgroupagents - 1
            ]
            [
              sprout 1 [set color black set shape "person" set g1opin 0.9 set g2opin -0.9 set group 1]
            ]
          ]
        ]
      ]
      set placedagents placedagents + 1
    ]
  ]
  set height height + 1
]
set height 0
set width width + 1
]
if majority-at-zero
[ask turtles [set g1opin 0]]
if authority-figures-on
[
  let loopy 0
  let change 0
  let target 0
  while [loopy < authority-figures]
  [
    set change 0
    while [change = 0]
    [
      set target random agent-numbers
      if not ([group] of turtle target = 3)
      [
        set change 1
      ]
    ]
    ask turtle target
    [
      set color red
      set g1opin 0.9
    ]
  ]
]

```

```

        set g2opin -0.9
        set group 3
    ]
    set loopy loopy + 1
]
set loopy 0
let dist 100000
while [loopy < count turtles]
[
    ifelse [group] of turtle loopy < 3
    [
        ask turtles with [group = 3] [if diff xcor [xcor] of turtle loopy + diff ycor [ycor] of turtle loopy < dist [set dist diff xcor [xcor] of turtle
loopy + diff ycor [ycor] of turtle loopy]]
        ifelse dist <= 3 [ask turtle loopy [set sanctioned-interactions 1]][ask turtle loopy [set sanctioned-interactions 0]]
        set dist 100000
    ]
    [
        ask turtle loopy [set sanctioned-interactions 0]
    ]
    set loopy loopy + 1
]
]
if common-interest
[
    let loopy 0
    while [loopy < agent-numbers]
    [
        ifelse random 2 = 1 [ask turtle loopy [set interest 1]][ask turtle loopy [set interest 2]]
        set loopy loopy + 1
    ]
]
]
end

to start-simulation
agents-interact
if unstigmatised-majority
[
    ifelse majority-at-zero
    [ask turtles [set g1opin 0]]
    [ask turtles [set g1opin 0.9]]
]
if mobility [move-agents]
draw-graph
set cycle cycle + 1
if cycle = 500 [create-file]
if cycle = 1000 [create-file]
if cycle = 3000 [create-file]
if cycle = cyclestorun
[
    create-file
    stop
]
; if cycle = 1000 ;;Probably needs to be 25000 to 30000 these days
; [
;     create-file
;     ifelse agents-in-minority < 120
;     [
;         set agents-in-minority agents-in-minority + 30
;     ]
;     [
;         set agents-in-minority 30
;         ifelse interaction-range < 24
;         [
;             set interaction-range interaction-range + 2
;         ]
;         [
;             set interaction-range 2
;             ifelse co-operation
;             [

```



```

let why [ycor] of turtle adjagent
set ecks ecks + 1
set alist lput find-turtle ecks why alist
if range > 3
[
  set why why + 1
  set alist lput find-turtle ecks why alist
]
if range > 5
[
  set ecks ecks - 1
  set alist lput find-turtle ecks why alist
]
if range > 7
[
  set ecks ecks - 1
  set alist lput find-turtle ecks why alist
]
if range > 9
[
  set ecks [xcor] of turtle adjagent + 2
  set why [ycor] of turtle adjagent
  set alist lput find-turtle ecks why alist
]
if range > 11
[
  set why why + 1
  set alist lput find-turtle ecks why alist
]
if range > 13
[
  set why why + 1
  set alist lput find-turtle ecks why alist
]
if range > 15
[
  set ecks ecks - 1
  set alist lput find-turtle ecks why alist
]
if range > 17
[
  set ecks ecks - 1
  set alist lput find-turtle ecks why alist
]
if range > 19
[
  set ecks ecks - 1
  set alist lput find-turtle ecks why alist
]
if range > 21
[
  set ecks ecks - 1
  set alist lput find-turtle ecks why alist
]
if range > 23
[
  set why why - 1
  set alist lput find-turtle ecks why alist
]
set alist remove adjagent alist
ifelse member? "Error: No turtle on patch searched for" alist
[print "Error: No turtle on patch searched for, patch listed above" report []]
[report alist]
end

to-report find-turtle [ecks why]
let limit dimension - 1
if ecks > limit [set ecks ecks - limit]
if why > limit [set why why - limit]
if ecks < 0 [set ecks ecks + limit]

```

```

if why < 0 [set why why + limit]
ifelse any? turtles-on patch ecks why
  [report [who] of one-of turtles-on patch ecks why]
  [write ecks print why report "Error: No turtle on patch searched for"]
end

;Note that for the two [group] = 3 if statements the original order was g2opin then g1opin then g1+g2/2
;This has been changed now!

to interact [agent1 agent2]
  let positive 0
  let a1a2 0
  let a2a1 0
  if [group] of turtle agent1 = 1
    [set a2a1 [g1opin] of turtle agent2]
  if [group] of turtle agent1 = 2
    [set a2a1 [g2opin] of turtle agent2]
  if [group] of turtle agent1 = 3
  [
    if [group] of turtle agent2 = 1
      [set a2a1 [g1opin] of turtle agent2]
    if [group] of turtle agent2 = 2
      [set a2a1 [g2opin] of turtle agent2]
    if [group] of turtle agent2 = 3
      [set a2a1 ([g1opin] of turtle agent2 + [g2opin] of turtle agent2) / 2]
  ]
  if [group] of turtle agent2 = 1
    [set a1a2 [g1opin] of turtle agent1]
  if [group] of turtle agent2 = 2
    [set a1a2 [g2opin] of turtle agent1]
  if [group] of turtle agent2 = 3
  [
    if [group] of turtle agent1 = 1
      [set a1a2 [g1opin] of turtle agent1]
    if [group] of turtle agent1 = 2
      [set a1a2 [g2opin] of turtle agent1]
    if [group] of turtle agent1 = 3
      [set a1a2 ([g1opin] of turtle agent1 + [g2opin] of turtle agent1) / 2]
  ]
  set positive a1a2 + a2a1
  if status-effect
  [
    if [group] of turtle agent1 = 1 and [group] of turtle agent2 = 2
    [
      set positive (a1a2 * status-effect-magnitude + a2a1 * (100 - status-effect-magnitude)) / 100
      set positive positive * 2 ;There are two opinions!
    ]
    if [group] of turtle agent1 = 2 and [group] of turtle agent2 = 1
    [
      set positive (a2a1 * status-effect-magnitude + a1a2 * (100 - status-effect-magnitude)) / 100
      set positive positive * 2 ;There are two opinions!
    ]
  ]
  set positive (positive + 2) * 25
  if common-interest
  [
    if random 100 < common-interest-perception and [interest] of turtle agent1 = [interest] of turtle agent2
    [
      let secondintpositive (positive - random 101) / 100
      change-opinions agent1 agent2 secondintpositive
      change-opinions agent2 agent1 secondintpositive
    ]
  ]
  if authority-figures-on
  [
    if ([group] of turtle agent1 = 1 and [group] of turtle agent2 = 2) or ([group] of turtle agent1 = 2 and [group] of turtle agent2 = 1)
    [
      if [sanctioned-interactions] of turtle agent1 = 1 and [sanctioned-interactions] of turtle agent2 = 1
      [
        set positive positive + 10

```

```

        let anotherintpositive (positive - random 101) / 100
        change-opinions agent1 agent2 anotherintpositive
        change-opinions agent2 agent1 anotherintpositive
    ]
]
]
set positive (positive - random 101) / 100 ;; Positivity of interaction in range -1 to 1
change-opinions agent1 agent2 positive
change-opinions agent2 agent1 positive
end

to change-opinions [agentnum secondagent positive]
let targetopinion [group] of turtle secondagent
let coop 0
if co-operation [if (random 1000) < co-op-freq * 10 [set coop 1]]
if targetopinion = 3
[
    if [group] of turtle agentnum = 1 [set targetopinion 1]
    if [group] of turtle agentnum = 2 [set targetopinion 2]
    if [group] of turtle agentnum = 3
    [
        ifelse random 2 < 1 [set targetopinion 1][set targetopinion 2]
    ]
]
let agentopinion 0
ifelse targetopinion = 1
[set agentopinion [g1opin] of turtle agentnum]
[set agentopinion [g2opin] of turtle agentnum]
let expectation agentopinion * 2
set expectation (expectation + 2) * 25
set expectation (expectation - 50) / 100
ifelse targetopinion = 1
[
    if expectation < positive
    [
        increase-opinion agentnum 1
        if coop = 1 [increase-opinion agentnum 1]
        set poscount poscount + 1
    ]
    if expectation > positive
    [
        decrease-opinion agentnum 1
        if coop = 1 [decrease-opinion agentnum 1]
        set negcount negcount + 1
    ]
]
[
    if expectation < positive
    [
        increase-opinion agentnum 2
        if coop = 1 [increase-opinion agentnum 2]
        set poscount poscount + 1
    ]
    if expectation > positive
    [
        decrease-opinion agentnum 2
        if coop = 1 [decrease-opinion agentnum 2]
        set negcount negcount + 1
    ]
]
clean-opinions agentnum
end

to increase-opinion [agentnum targetgroup]
ifelse targetgroup = 1
[
    ifelse learning-rule
    [
        ask turtle agentnum [set g1opin ([g1opin] of turtle agentnum + 0.1) - ([g1opin] of turtle agentnum * 0.1)]
    ]
]

```

```

[
  ask turtle agentnum [set g1opin [g1opin] of turtle agentnum + 0.0001]
  if not no-limits [if [g1opin] of turtle agentnum > 1 [ask turtle agentnum [set g1opin 1]]]
]
]
[
  ifelse learning-rule
  [
    ask turtle agentnum [set g1opin ((g2opin] of turtle agentnum + 0.1) - ((g2opin] of turtle agentnum * 0.1)]
  ]
  [
    ask turtle agentnum [set g2opin [g2opin] of turtle agentnum + 0.0001]
    if not no-limits [if [g2opin] of turtle agentnum > 1 [ask turtle agentnum [set g2opin 1]]]
  ]
]
]
end

```

```

to decrease-opinion [agentnum targetgroup]
  ifelse targetgroup = 1
  [
    ifelse learning-rule
    [
      ask turtle agentnum [set g1opin ((g1opin] of turtle agentnum - 0.1) - ((g1opin] of turtle agentnum * 0.1)]
    ]
    [
      ask turtle agentnum [set g1opin [g1opin] of turtle agentnum - 0.0001]
      if not no-limits [if [g1opin] of turtle agentnum < -1 [ask turtle agentnum [set g1opin -1]]]
    ]
  ]
  [
    ifelse learning-rule
    [
      ask turtle agentnum [set g1opin ((g2opin] of turtle agentnum - 0.1) - ((g2opin] of turtle agentnum * 0.1)]
    ]
    [
      ask turtle agentnum [set g2opin [g2opin] of turtle agentnum - 0.0001]
      if not no-limits [if [g2opin] of turtle agentnum < -1 [ask turtle agentnum [set g2opin -1]]]
    ]
  ]
]
end

```

```

to-report difference [value1 value2]
  ifelse value1 > value2
  [report value1 - value2]
  [report value2 - value1]
end

```

```

to clean-opinions [agentnum]
  let g1 [g1opin] of turtle agentnum
  let g2 [g2opin] of turtle agentnum
  set g1 g1 * 10000
  set g1 round g1
  set g1 g1 / 10000
  set g2 g2 * 10000
  set g2 round g2
  set g2 g2 / 10000
  ask turtle agentnum [set g1opin g1]
  ask turtle agentnum [set g2opin g2]
end

```

```

to move-agents
  let a1 random agent-numbers
  let a2 random agent-numbers
  let xstore [xcor] of turtle a1
  let ystore [ycor] of turtle a1
  ask turtle a1
  [
    set xcor [xcor] of turtle a2
    set ycor [ycor] of turtle a2
  ]
]

```

```

ask turtle a2
[
  set xcor xstore
  set ycor ystore
]
end

to draw-graph
set-current-plot "Opinions"
set-current-plot-pen "G1G1"
plot mean [g1opin] of turtles with [group = 1]
set-current-plot-pen "G1G2"
plot mean [g2opin] of turtles with [group = 1]
set-current-plot-pen "G2G1"
plot mean [g1opin] of turtles with [group = 2]
ifelse PrevSim
[
  set-current-plot-pen "G1G2Prev"
  plot item prevplot prevvalues
  set prevplot prevplot + 1
]
[
  set prevvalues lput mean [g2opin] of turtles with [group = 1] prevvalues
]
set-current-plot-pen "G2G2"
plot mean [g2opin] of turtles with [group = 2]
set-current-plot-pen "zero"
plot 0
; ask turtles
; [
;   set color rgb 0 ((g2opin * 100) + 100) 0
; ]
end

to create-file
let loopy 0
let authorities []
while [loopy < agent-numbers]
[
  ifelse authority-figures-on
  [
    let dist 100000
    ask turtles with [group = 3] [if diff xcor [xcor] of turtle loopy + diff ycor [ycor] of turtle loopy < dist [set dist diff xcor [xcor] of turtle
loopy + diff ycor [ycor] of turtle loopy]]
    set authorities lput dist authorities
  ]
  [
    set authorities lput -999 authorities
  ]
  set loopy loopy + 1
]
let int-range interaction-range
let co-op 0
if co-operation [set co-op co-op-freq]
let status 50
if status-effect [set status status-effect-magnitude]
let filename "Results500.txt"
if file-exists? filename [set filename "Results1000.txt"]
if file-exists? filename [set filename "Results3000.txt"]
if file-exists? filename [set filename "Results5000.txt"]
file-open filename
file-type (word "AvgG2opin MinSize AuthorityAgents IntRange Coop CommonInterests Status" "\n")
file-close
set loopy 0
let agents 0
let totalopin 0
while [loopy < agent-numbers]
[
  if [group] of turtle loopy = 1
  [

```

```

    set agents agents + 1
    set totalopin totalopin + [g2opin] of turtle loopy
  ]
  set loopy loopy + 1
]
set totalopin totalopin / agents
set totalopin precision totalopin 4
let line (word totalopin " " agents-in-minority " " authority-figures " " int-range " " co-op " " common-interest-perception " " status)
file-open (word filename)
file-type (word line "\n")
file-close
end

```

```

to-report diff [ecks why]
  ifelse ecks > why [report ecks - why][report why - ecks]
end

```

DGM

```

globals [learning-rate cycle moves initagents]
turtles-own [id segregation mates happy]

```

```

to setup-model
  ;; (for this model to work with NetLogo's new plotting features,
  ;; __clear-all-and-reset-ticks should be replaced with clear-all at
  ;; the beginning of your setup procedure and reset-ticks at the end
  ;; of the procedure.)
  __clear-all-and-reset-ticks
  set cycle 0
  set learning-rate 0.01
  let population round ((population-density / 100) * count patches)
  ask n-of population patches
  [
    sprout 1 [ set color black set shape "circle" set id 0 set happy 1 set color approximate-rgb 50 50 50]
  ]
  let groupsize round (0.5 * count turtles with [xcor < 20 and ycor < 20])
  ask n-of groupsize turtles with [xcor < 20 and ycor < 20]
  [
    set id learning-rate
    set color approximate-rgb 0 0 (id * 255)
  ]
  set initagents count turtles with [id > 0]
end

```

```

to start-simulation
  tick
  set moves 0
  set cycle cycle + 1
  update-colours
  repeat int-freq
  [
    agents-interact
  ]
  if move-freq = 0
  [
    ;; We don't really care about happiness, but it calculates a segregation statistic that we may still need
    check-happiness
  ]
  repeat move-freq
  [
    check-happiness
    move-turtles
  ]
  if cycle = 500
  [
    save-state
    next-simulation
  ]
end

```

```

to update-colours
  ask turtles with [id > 0]
  [
    set color approximate-rgb 0 0 (100 + id * 155)
  ]
  ask turtles with [id = 0]
  [
    set color approximate-rgb 50 50 50
  ]
end

to agents-interact
  let loopy 0
  let outgroupinteractions 0
  let ingroupinteractions 0
  while [loopy < count turtles]
  [
    if count turtles-on [neighbors] of turtle loopy > 0
    [
      let adjacents turtles [who] of turtles-on [neighbors] of turtle loopy
      foreach adjacents
      [
        interact loopy ?
      ]
    ]
    set loopy loopy + 1
  ]
end

to interact [turtle1 turtle2]
  let g1 [id] of turtle turtle1
  let g2 [id] of turtle turtle2
  ifelse g1 > 0 [set g1 1][set g1 0]
  ifelse g2 > 0 [set g2 1][set g2 0]
  ifelse g1 = 1 and g2 = 1
  [
    ingroup-interaction turtle1 turtle2
  ]
  [
    ifelse g1 = 0 and g2 = 0
    [
      ;;No point doing interactions between two outgroup members
    ]
    [
      outgroup-interaction turtle1 turtle2
    ]
  ]
end

to ingroup-interaction [t1 t2]
  ifelse random(100) <= blue-pos-interactions
  [
    increase-id t1
    increase-id t2
  ]
  [
    decrease-id t1
    decrease-id t2
  ]
end

to outgroup-interaction [t1 t2]
  if [id] of turtle t2 > 0
  [set t1 t2]
  increase-id t1
end

to increase-id [t1]
  ask turtle t1 [set id [id] of turtle t1 + learning-rate]
  if [id] of turtle t1 > 1 [ask turtle t1 [set id 1]]

```

```

end

to decrease-id [t1]
  ask turtle t1 [set id [id] of turtle t1 - learning-rate]
  if [id] of turtle t1 < 0 [ask turtle t1 [set id 0]]
end

;;;START CODE FOR SCHELLINGS STUFF!!
*****
*****
;;;START CODE FOR SCHELLINGS STUFF!!
*****
*****
;;;START CODE FOR SCHELLINGS STUFF!!
*****
*****

to check-happiness
  let loopy 0
  let same 0
  let different 0
  while [loopy < count turtles]
  [
    if [id] of turtle loopy > 0
    [
      set same count (turtles-on [neighbors] of turtle loopy) with [id > 0]
      ask turtle loopy [set mates same]
      set different count (turtles-on [neighbors] of turtle loopy) with [id = 0]
      ifelse same + different = 0
      [
        ask turtle loopy [set happy 1]
        set same 1
        set different 0
      ]
      [
        ifelse (same / (same + different)) * 100 >= happy-percent
        [ask turtle loopy [set happy 1]]
        [ask turtle loopy [set happy 0]]
      ]
    ]
  ]
  if [id] of turtle loopy = 0
  [
    ask turtle loopy [set happy 1]
    set same count (turtles-on [neighbors] of turtle loopy) with [id = 0]
    set different count (turtles-on [neighbors] of turtle loopy) with [id > 0]
    if different > same
    [
      if random 100 = 1
      [
        increase-id loopy
      ]
    ]
    ifelse same + different = 0
    [
      ask turtle loopy [set happy 1]
      set same 1
      set different 0
    ]
    [
      ifelse (same / (same + different)) * 100 >= happy-percent
      [ask turtle loopy [set happy 1]]
      [ask turtle loopy [set happy 0]]
    ]
  ]
  ask turtle loopy [set segregation same / (same + different)]
  set loopy loopy + 1
]
end

```



```

to move-turtles
  let loopy 0
  while [loopy < count turtles]
  [
    if [happy] of turtle loopy = 0
    [
      set moves moves + 1
      ask turtle loopy [move-to one-of patches with [count turtles-here = 0]]
    ]
    set loopy loopy + 1
  ]
end

;;;START CODE FOR SAVING STUFF!!
*****
;;;START CODE FOR SAVING STUFF!!
*****
;;;START CODE FOR SAVING STUFF!!
*****

to save-grid
  file-open "Grid.txt"
  let x 0
  let y 0
  while [x < world-width]
  [
    while [y < world-height]
    [
      ifelse count turtles-on patch x y > 0
      [
        if count turtles-on patch x y > 1
        [
          file-print "ERROR 2 turtles on a patch!"
        ]
        ifelse [id] of one-of turtles-on patch x y > 0
        [
          file-type "1"
        ]
        [
          file-type "0"
        ]
      ]
      [
        file-type " "
      ]
    ]
    set y y + 1
  ]
  set y 0
  file-print ""
  set x x + 1
]
file-close
end

to save-state
  file-open "S21Result.txt"
  ;;;File format: Happy percent, Ingroup pos interactions, inital size, average identification, group size, seg1, seg2, mates, moves, outsiders
  file-type happy-percent
  file-type " "
  file-type blue-pos-interactions
  file-type " "
  file-type initagents
  file-type " "
  file-type mean [id] of turtles with [id > 0]
  file-type " "
  file-type count turtles with [id > 0]
  file-type " "

```

```

file-type mean [segregation] of turtles with [id > 0]
file-type " "
file-type mean [segregation] of turtles with [id > 0] / count turtles with [id > 0] * 1000
file-type " "
file-type mean [mates] of turtles with [id > 0]
file-type " "
file-type moves
file-type " "
file-print count turtles with [(xcor > 20 or ycor > 20) and id > 0]
file-close
end

;;;START CODE FOR MULTIPLE SIMULATIONS!!
*****
**
;;;START CODE FOR MULTIPLE SIMULATIONS!!
*****
**
;;;START CODE FOR MULTIPLE SIMULATIONS!!
*****
**

to next-simulation
  ifelse replication < 10
  [
    set replication replication + 1
    setup-model
  ]
  [
    set replication 1
    ifelse happy-percent < 100
    [
      set happy-percent happy-percent + 10
      setup-model
    ]
    [
      ifelse blue-pos-interactions < 100
      [
        set happy-percent 0
        set blue-pos-interactions blue-pos-interactions + 10
        setup-model
      ]
      [
        stop
      ]
    ]
  ]
end

```

E-DGM

```

globals [learning-rate cycle moves initagents]
turtles-own [id1 id2 id3 id4 segregation mates happy]

to setup-model
  ;; (for this model to work with NetLogo's new plotting features,
  ;; __clear-all-and-reset-ticks should be replaced with clear-all at
  ;; the beginning of your setup procedure and reset-ticks at the end
  ;; of the procedure.)
  __clear-all-and-reset-ticks
  set cycle 0
  set learning-rate 0.01
  let population round ((population-density / 100) * count patches)
  set population (round (population / 4)) * 4 ;This moves the population to the closet number divisable by four
  ;Start by placing the right number of turtles
  ask n-of population patches
  [
    sprout 1 [ set color black set shape "circle" set happy 1 set color approximate-rgb 30 30 30]
  ]
  ;Then make half of the turtles in each region join that regions group

```

```

ask n-of (round ((population-density / 100) * 312.5)) turtles with [xcor < 20 and ycor < 20]
[
  set id1 0.01
]
ask n-of (round ((population-density / 100) * 312.5)) turtles with [xcor < 20 and ycor >= 25 and ycor <= 45]
[
  set id2 0.01
]
ask n-of (round ((population-density / 100) * 312.5)) turtles with [xcor >= 25 and xcor <= 45 and ycor < 20]
[
  set id3 0.01
]
ask n-of (round ((population-density / 100) * 312.5)) turtles with [xcor >= 25 and xcor <= 45 and ycor >= 25 and ycor <= 45]
[
  set id4 0.01
]
;Now all of the turtles have one identity, the last step is to introduce multiple identities DOES NOT WORKY!!!
let midturtles round ((mult-ids / 100) * count turtles)
set midturtles round (midturtles / 12) ;This gives how many multiple identity turtles of each type are recovered. Now to go through all
twelve conditions...
ask n-of midturtles turtles with [id1 = 0.5]
[
  set id2 0.5
]
ask n-of midturtles turtles with [id1 = 0.5]
[
  set id3 0.5
]
ask n-of midturtles turtles with [id1 = 0.5]
[
  set id4 0.5
]
ask n-of midturtles turtles with [id2 = 0.5 and id1 = 0]
[
  set id1 0.5
]
ask n-of midturtles turtles with [id2 = 0.5 and id1 = 0]
[
  set id3 0.5
]
ask n-of midturtles turtles with [id2 = 0.5 and id1 = 0]
[
  set id4 0.5
]
ask n-of midturtles turtles with [id3 = 0.5 and id1 = 0 and id2 = 0]
[
  set id1 0.5
]
ask n-of midturtles turtles with [id3 = 0.5 and id1 = 0 and id2 = 0]
[
  set id2 0.5
]
ask n-of midturtles turtles with [id3 = 0.5 and id1 = 0 and id2 = 0]
[
  set id4 0.5
]
ask n-of midturtles turtles with [id4 = 0.5 and id1 = 0 and id2 = 0 and id3 = 0]
[
  set id1 0.5
]
ask n-of midturtles turtles with [id4 = 0.5 and id1 = 0 and id2 = 0 and id3 = 0]
[
  set id2 0.5
]
ask n-of midturtles turtles with [id4 = 0.5 and id1 = 0 and id2 = 0 and id3 = 0]
[
  set id3 0.5
]
]
;All of the turtles should now be ready to go. Just need to update the colours to make it visually apparent what's going on
update-color

```

```

end

to update-color
  let redc 0
  let greenc 0
  let bluec 0
  let whitec 0
  let high -1
  let low 9999
  let idcount 0
  ask turtles
  [
    set redc 20 + id4 * 90 + id1 * 135
    set greenc 20 + id4 * 90 + id2 * 135
    set bluec 20 + id4 * 90 + id3 * 135
    set color approximate-rgb redc greenc bluec
    set idcount 0
    if id1 > 0 [set idcount idcount + 1]
    if id2 > 0 [set idcount idcount + 1]
    if id3 > 0 [set idcount idcount + 1]
    if id4 > 0 [set idcount idcount + 1]
    if idcount > 1 [set shape "circle"]
    if idcount <= 1 [set shape "square"]
  ]
end

to start-simulation
  tick
  set moves 0
  set cycle cycle + 1
  repeat int-freq [agents-interact]
  check-happiness
  repeat move-freq [move-turtles]
  agents-conform ;Used to be part of the happiness check, moved to its own function for readability

; if cycle = 1 [export-view "High1.png"]
; if cycle = 50 [export-view "High50.png"]
; if cycle = 100 [export-view "High100.png"]
; if cycle = 150 [export-view "High150.png"]
; if cycle = 200 [export-view "High200.png"]
; if cycle = 250 [export-view "High250.png"]
; if cycle = 300 [export-view "High300.png"]
; if cycle = 350 [export-view "High350.png"]
; if cycle = 400 [export-view "High400.png"]
; if cycle = 450 [export-view "High450.png"]
; if cycle = 500 [export-view "High500.png" stop]

if cycle = 2000
[
; stop
; save-state
; next-simulation
]
if update-colors [update-color]
end

to agents-interact
  let loopy 0
  let outgroupinteractions 0 ;These give stats for the round
  let ingroupinteractions 0 ;Ditto
  while [loopy < count turtles]
  [
    let ecks [xcor] of turtle loopy
    let why [ycor] of turtle loopy
    set ecks ecks + 1
    if ecks > 49 [set ecks 0]
    if count turtles-on patch ecks why = 1 [interact [who] of one-of turtles-on patch ecks why loopy]
    set why why - 1
  ]

```

```

    if why < 0 [set why 49]
    if count turtles-on patch ecks why = 1 [interact [who] of one-of turtles-on patch ecks why loopy]
    set ecks ecks - 1
    if ecks < 0 [set ecks 49]
    if count turtles-on patch ecks why = 1 [interact [who] of one-of turtles-on patch ecks why loopy]
    set ecks ecks - 1
    if ecks < 0 [set ecks 49]
    if count turtles-on patch ecks why = 1 [interact [who] of one-of turtles-on patch ecks why loopy]
    set loopy loopy + 1
  ]
end

to interact [t1 t2]
  let ingroup 0
  let sharedgroups []
  ;Start by establishing whether we've got an ingroup or outgroup interaction
  if [id1] of turtle t1 > 0 and [id1] of turtle t2 > 0 [set ingroup 1 set sharedgroups lput 1 sharedgroups]
  if [id2] of turtle t1 > 0 and [id2] of turtle t2 > 0 [set ingroup 1 set sharedgroups lput 2 sharedgroups]
  if [id3] of turtle t1 > 0 and [id3] of turtle t2 > 0 [set ingroup 1 set sharedgroups lput 3 sharedgroups]
  if [id4] of turtle t1 > 0 and [id4] of turtle t2 > 0 [set ingroup 1 set sharedgroups lput 4 sharedgroups]
  ;Now find the relevant group
  let t1rel []
  let t2rel []
  ifelse ingroup = 1
  [
    set t1rel one-of sharedgroups
    set t2rel t1rel
  ]
  [
    let t1list []
    let t2list []
    if [id1] of turtle t1 > 0 [set t1list lput 1 t1list]
    if [id2] of turtle t1 > 0 [set t1list lput 2 t1list]
    if [id3] of turtle t1 > 0 [set t1list lput 3 t1list]
    if [id4] of turtle t1 > 0 [set t1list lput 4 t1list]
    if [id1] of turtle t2 > 0 [set t2list lput 1 t2list]
    if [id2] of turtle t2 > 0 [set t2list lput 2 t2list]
    if [id3] of turtle t2 > 0 [set t2list lput 3 t2list]
    if [id4] of turtle t2 > 0 [set t2list lput 4 t2list]
    ifelse length t1list > 0
    [set t1rel one-of t1list]
    [set t1rel 5] ;Dummy number, should change nothing
    ifelse length t2list > 0
    [set t2rel one-of t2list]
    [set t2rel 5] ;Dummy number, should change nothing
  ]
  ;Now do the identity changes to the relevant groups
  ifelse ingroup = 1
  [
    ingroup-interaction t1 t2 t1rel t2rel
  ]
  [
    outgroup-interaction t1 t2 t1rel t2rel
  ]
end

to ingroup-interaction [t1 t2 i1 i2]
  ;Ingroup interactions can lead to an increase or decrease
  ifelse random(100) < pos-interactions
  [
    increase-id t1 i1
    increase-id t2 i2
  ]
  [
    decrease-id t1 i1
    decrease-id t2 i2
  ]
end

to outgroup-interaction [t1 t2 i1 i2]

```

```

;Outgroup interactions always lead to an identity increase
increase-id t1 i1
increase-id t2 i2
end

to increase-id [agent identity]
  if identity = 1 [ask turtle agent [set id1 ([id1] of turtle agent + learning-rate) - ([id1] of turtle agent * learning-rate)]]
  if identity = 2 [ask turtle agent [set id2 ([id2] of turtle agent + learning-rate) - ([id2] of turtle agent * learning-rate)]]
  if identity = 3 [ask turtle agent [set id3 ([id3] of turtle agent + learning-rate) - ([id3] of turtle agent * learning-rate)]]
  if identity = 4 [ask turtle agent [set id4 ([id4] of turtle agent + learning-rate) - ([id4] of turtle agent * learning-rate)]]
end

to decrease-id [agent identity]
  if identity = 1 [ask turtle agent [set id1 ([id1] of turtle agent - learning-rate) - ([id1] of turtle agent * learning-rate)]]
  if identity = 2 [ask turtle agent [set id2 ([id2] of turtle agent - learning-rate) - ([id2] of turtle agent * learning-rate)]]
  if identity = 3 [ask turtle agent [set id3 ([id3] of turtle agent - learning-rate) - ([id3] of turtle agent * learning-rate)]]
  if identity = 4 [ask turtle agent [set id4 ([id4] of turtle agent - learning-rate) - ([id4] of turtle agent * learning-rate)]]
  if [id1] of turtle agent < 0 [ask turtle agent [set id1 0]]
  if [id2] of turtle agent < 0 [ask turtle agent [set id2 0]]
  if [id3] of turtle agent < 0 [ask turtle agent [set id3 0]]
  if [id4] of turtle agent < 0 [ask turtle agent [set id4 0]]
end

to check-happiness
  let same 0
  let different 0
  let thisone 0
  let loopy 0
  let neighbours []
  ask turtles
  [
    set same 0
    set different 0
    set neighbours turtles-on neighbors
    set neighbours [who] of neighbours
    set loopy 0
    while [loopy < length neighbours] ;This hellish bit is iterating through all neighbours and checking to see if they share a group or not. It's
    a bottleneck.
    [
      set thisone 0
      if id1 > 0 and [id1] of turtle item loopy neighbours > 0 [set thisone 1]
      if id2 > 0 and [id2] of turtle item loopy neighbours > 0 [set thisone 1]
      if id3 > 0 and [id3] of turtle item loopy neighbours > 0 [set thisone 1]
      if id4 > 0 and [id4] of turtle item loopy neighbours > 0 [set thisone 1]
      ifelse thisone = 1 [set same same + 1][set different different + 1]
      set loopy loopy + 1
    ]
    ifelse same + different = 0
    [
      set happy 1
      set same 1
      set different 0
    ]
    [
      ifelse (same / (same + different)) * 100 >= similarity-preference
      [set happy 1]
      [set happy 0]
    ]
    set segregation same / (same + different)
  ]
end

to move-turtles
  ask turtles
  [
    if happy = 0
    [
      set moves moves + 1
      move-to one-of patches with [count turtles-here = 0]
    ]
  ]
end

```

```

]
end

to agents-conform
ask turtles
[
  let neighbours turtles-on neighbors
  let i1a count neighbours with [id1 > 0]
  let i2a count neighbours with [id2 > 0]
  let i3a count neighbours with [id3 > 0]
  let i4a count neighbours with [id4 > 0]
  let highball i1a
  let conformto 1
  if i2a > highball [set highball i2a set conformto 2]
  if i3a > highball [set highball i3a set conformto 3]
  if i4a > highball [set highball i4a set conformto 4]
  if highball > count neighbours / 2 ;i.e. if more than half the adjacent agents have this identity
  [
    if random(100) = 0
    [
      ;Conforming time!
      if conformto = 1 [if id1 = 0 [set id1 learning-rate]] ;Setting to learning rate is mathematically equivalent to an increase-id operation
      if conformto = 2 [if id2 = 0 [set id2 learning-rate]]
      if conformto = 3 [if id3 = 0 [set id3 learning-rate]]
      if conformto = 4 [if id4 = 0 [set id4 learning-rate]]
    ]
  ]
]
end

to save-state
file-open "RanTest.txt"
;;File format: Similarity Preference, Ingroup pos interactions, initial size, average identification, group size, seg1, seg2, mates, moves,
outsiders, g2size, g3size, g4size, multids
file-type similarity-preference
file-type " "
file-type pos-interactions
file-type " "
file-type initagents
file-type " "
file-type mean [id1] of turtles with [id1 > 0]
file-type " "
file-type count turtles with [id1 > 0]
file-type " "
file-type mean [segregation] of turtles with [id1 > 0]
file-type " "
file-type mean [segregation] of turtles with [id1 > 0] / count turtles with [id1 > 0] * 1000
file-type " "
file-type mean [mates] of turtles with [id1 > 0]
file-type " "
file-type moves
file-type " "
file-type count turtles with [(xcor > 24 or ycor > 24) and id1 > 0]
file-type " "
file-type count turtles with [id2 > 0]
file-type " "
file-type count turtles with [id3 > 0]
file-type " "
file-type count turtles with [id4 > 0]
file-type " "
let loopy 0
let ids 0
let mids 0
while [loopy < count turtles]
[
  set ids 0
  if [id1] of turtle loopy > 0 [set ids ids + 1]
  if [id2] of turtle loopy > 0 [set ids ids + 1]
  if [id3] of turtle loopy > 0 [set ids ids + 1]
  if [id4] of turtle loopy > 0 [set ids ids + 1]
]

```

```

    if ids > 1 [set mids mids + 1]
    set loopy loopy + 1
  ]
  file-print mids
  file-close
end

to next-simulation
  ifelse replication < 10
  [
    set replication replication + 1
    setup-model
  ]
  [
    set replication 1
    ifelse similarity-preference < 100
    [
      set similarity-preference similarity-preference + 10
      setup-model
    ]
    [
      ifelse pos-interactions < 100
      [
        set similarity-preference 0
        set pos-interactions pos-interactions + 10
        setup-model
      ]
      [
        stop
      ]
    ]
  ]
]
end

to save-grid
  ;Save grid function not implemented at the moment (clustering rules)
end

;to save-grid
; file-open "Grid.txt"
; let x 0
; let y 0
; while [x < world-width]
; [
;   while [y < world-height]
;   [
;     ifelse count turtles-on patch x y > 0
;     [
;       if count turtles-on patch x y > 1
;       [
;         file-print "ERROR 2 turtles on a patch!"
;       ]
;     ]
;     ifelse [id] of one-of turtles-on patch x y > 0
;     [
;       file-type "1"
;     ]
;     [
;       file-type "0"
;     ]
;   ]
;   [
;     file-type " "
;   ]
;   set y y + 1
; ]
; set y 0
; file-print ""
; set x x + 1
; ]
; file-close

```



```
;end
```

MGSM

```
globals [segregation redseg blueseg greenseg cycle chaos]
turtles-own [happy native]

to setup-model
  ;; (for this model to work with NetLogo's new plotting features,
  ;; __clear-all-and-reset-ticks should be replaced with clear-all at
  ;; the beginning of your setup procedure and reset-ticks at the end
  ;; of the procedure.)
  __clear-all-and-reset-ticks
  if multtid and two-groups
  [
    print "That's a stupid parameter setup, doing put multi id agents and twogroups together"
    stop
  ]
  set cycle 1
  set chaos 0
  ask n-of (count patches * 5 / 6) patches
  [
    sprout 1 [ set color red set happy 1 ]
  ]
  ask n-of (count turtles / 2) turtles [ set color green ]
  if not two-groups [ask n-of (count turtles / 3) turtles [ set color blue ]]
  if multtid
  [
    let multidturtles round (count turtles / 100 * multtid-percent)
    let rgnum round ((rg-ratio / (rg-ratio + rb-ratio + bg-ratio)) * multidturtles)
    let rbnum round ((rb-ratio / (rg-ratio + rb-ratio + bg-ratio)) * multidturtles)
    let bgnum round ((bg-ratio / (rg-ratio + rb-ratio + bg-ratio)) * multidturtles)
    set multidturtles rgnum + rbnum + bgnum
    let idset n-of multidturtles turtles
    ask n-of rgnum idset with [shape = "default"]
    [
      set shape "circle"
      set color yellow
    ]
    ask n-of rbnum idset with [shape = "default"]
    [
      set shape "circle"
      set color magenta
    ]
    ask n-of bgnum idset with [shape = "default"]
    [
      set shape "circle"
      set color cyan
    ]
  ]
]
end

to start-simulation
  check-happiness
  move-turtles
; update-stats ;;Remember to comment this out to make the thing run faster!
  set cycle cycle + 1
  if count turtles with [happy = 0] = 0
  [
    if become-native [nativify]
    update-stats
    ;next-simulation
  ]
  if cycle > 200
  [
    set chaos 1
    if become-native [nativify]
    update-stats
  ]
end
```

```

;next-simulation
]
end

to check-happiness
let loopy 0
let same 0
let different 0
while [loopy < count turtles]
[
if [color] of turtle loopy = red
[
set same count (turtles-on [neighbors] of turtle loopy) with [color = red or color = yellow or color = magenta]
set different count (turtles-on [neighbors] of turtle loopy) with [color = green or color = blue or color = cyan]
ifelse same + different > 0
[
ifelse (same / (same + different)) * 100 >= happy-percent
[ask turtle loopy [set happy 1]]
[ask turtle loopy [set happy 0]]
]
[
ask turtle loopy [set happy 0]
]
]
]
if [color] of turtle loopy = green
[
set same count (turtles-on [neighbors] of turtle loopy) with [color = green or color = yellow or color = cyan]
set different count (turtles-on [neighbors] of turtle loopy) with [color = red or color = blue or color = magenta]
ifelse same + different > 0
[
ifelse (same / (same + different)) * 100 >= happy-percent
[ask turtle loopy [set happy 1]]
[ask turtle loopy [set happy 0]]
]
[
ask turtle loopy [set happy 0]
]
]
]
if [color] of turtle loopy = blue
[
set same count (turtles-on [neighbors] of turtle loopy) with [color = blue or color = magenta or color = cyan]
set different count (turtles-on [neighbors] of turtle loopy) with [color = red or color = green or color = yellow]
ifelse same + different > 0
[
ifelse (same / (same + different)) * 100 >= happy-percent
[ask turtle loopy [set happy 1]]
[ask turtle loopy [set happy 0]]
]
[
ask turtle loopy [set happy 0]
]
]
]
if [color] of turtle loopy = cyan
[
set same count (turtles-on [neighbors] of turtle loopy) with [color = green or color = blue or color = cyan or color = magenta or color =
yellow]
set different count (turtles-on [neighbors] of turtle loopy) with [color = red]
ifelse same + different > 0
[
ifelse (same / (same + different)) * 100 >= happy-percent
[ask turtle loopy [set happy 1]]
[ask turtle loopy [set happy 0]]
]
[
ask turtle loopy [set happy 0]
]
]
]
if [color] of turtle loopy = magenta
[
set same count (turtles-on [neighbors] of turtle loopy) with [color = red or color = blue or color = cyan or color = magenta or color =

```

```

yellow]
  set different count (turtles-on [neighbors] of turtle loopy) with [color = green]
  ifelse same + different > 0
  [
    ifelse (same / (same + different)) * 100 >= happy-percent
    [ask turtle loopy [set happy 1]]
    [ask turtle loopy [set happy 0]]
  ]
  [
    ask turtle loopy [set happy 0]
  ]
]
if [color] of turtle loopy = yellow
[
  set same count (turtles-on [neighbors] of turtle loopy) with [color = green or color = red or color = cyan or color = magenta or color =
yellow]
  set different count (turtles-on [neighbors] of turtle loopy) with [color = blue]
  ifelse same + different > 0
  [
    ifelse (same / (same + different)) * 100 >= happy-percent
    [ask turtle loopy [set happy 1]]
    [ask turtle loopy [set happy 0]]
  ]
  [
    ask turtle loopy [set happy 0]
  ]
]
set loopy loopy + 1
]
end

to move-turtles
  let loopy 0
  while [loopy < count turtles]
  [
    if [happy] of turtle loopy = 0
    [
      ask turtle loopy [move-to one-of patches with [count turtles-here = 0]]
    ]
    set loopy loopy + 1
  ]
end

to swap-turtles [t1 t2]
  let tempx [xcor] of turtle t1
  let tempy [ycor] of turtle t1
  ask turtle t1 [set xcor [xcor] of turtle t2]
  ask turtle t1 [set ycor [ycor] of turtle t2]
  ask turtle t2 [set xcor tempx]
  ask turtle t2 [set ycor tempy]
end

to update-stats
  let loopy 0
  let workingseg 0
  let wredseg 0
  let wgreenseg 0
  let wblueseg 0
  let same 0
  let different 0
  while [loopy < count turtles]
  [
    if [color] of turtle loopy = red
    [
      set same count (turtles-on [neighbors] of turtle loopy) with [color = red or color = yellow or color = magenta]
      set different count (turtles-on [neighbors] of turtle loopy) with [color = green or color = blue or color = cyan]
      ifelse same + different > 0
      [set wredseg wredseg + (same / (same + different))]
      [set wredseg wredseg + 1]
    ]
  ]
]

```

```

if [color] of turtle loopy = green
[
  set same count (turtles-on [neighbors] of turtle loopy) with [color = green or color = yellow = color = cyan]
  set different count (turtles-on [neighbors] of turtle loopy) with [color = red or color = blue or color = magenta]
  ifelse same + different > 0
    [set wgreenseg wgreenseg + (same / (same + different))]
    [set wgreenseg wgreenseg + 1]
]
if [color] of turtle loopy = blue
[
  set same count (turtles-on [neighbors] of turtle loopy) with [color = blue or color = magenta or color = cyan]
  set different count (turtles-on [neighbors] of turtle loopy) with [color = red or color = green or color = yellow]
  ifelse same + different > 0
    [set wblueseg wblueseg + (same / (same + different))]
    [set wblueseg wblueseg + 1]
]
if [color] of turtle loopy = cyan
[
  set same count (turtles-on [neighbors] of turtle loopy) with [color = green or color = blue or color = cyan or color = magenta or color =
yellow]
  set different count (turtles-on [neighbors] of turtle loopy) with [color = red]
]
if [color] of turtle loopy = magenta
[
  set same count (turtles-on [neighbors] of turtle loopy) with [color = red or color = blue or color = cyan or color = magenta or color =
yellow]
  set different count (turtles-on [neighbors] of turtle loopy) with [color = green]
]
if [color] of turtle loopy = yellow
[
  set same count (turtles-on [neighbors] of turtle loopy) with [color = red or color = green or color = cyan or color = magenta or color =
yellow]
  set different count (turtles-on [neighbors] of turtle loopy) with [color = blue]
]
ifelse same + different > 0
  [set workingseg workingseg + (same / (same + different))]
  [set workingseg workingseg + 1]
set loopy loopy + 1
]
set workingseg (workingseg / count turtles) * 100
set wredseg (wredseg / count turtles with [color = red]) * 100
set wgreenseg (wgreenseg / count turtles with [color = green]) * 100
if not two-groups [set wblueseg (wblueseg / count turtles with [color = blue]) * 100]
set segregation workingseg
set redseg wredseg
set greenseg wgreenseg
set blueseg wblueseg
end

to nativify
let loopy 0
while [loopy < count turtles]
[
  if [shape] of turtle loopy = "circle"
  [
    if [color] of turtle loopy = yellow [yellownative loopy]
    if [color] of turtle loopy = cyan [cyannative loopy]
    if [color] of turtle loopy = magenta [magentanative loopy]
  ]
  set loopy loopy + 1
]
ask turtles with [shape = "circle"]
[
  set color native
]
end

to yellownative [turtlenum]
let redcount count (turtles-on [neighbors] of turtle turtlenum) with [color = red]
let greencount count (turtles-on [neighbors] of turtle turtlenum) with [color = green]

```

```

if redcount > greencount [ask turtle turtlenum [set native red]]
if redcount < greencount [ask turtle turtlenum [set native green]]
if redcount = greencount [ifelse random 2 = 1 [ask turtle turtlenum [set native red]][ask turtle turtlenum [set native green]]]
end

to cyannative [turtlenum]
let bluecount count (turtles-on [neighbors] of turtle turtlenum) with [color = blue]
let greencount count (turtles-on [neighbors] of turtle turtlenum) with [color = green]
if bluecount > greencount [ask turtle turtlenum [set native blue]]
if bluecount < greencount [ask turtle turtlenum [set native green]]
if bluecount = greencount [ifelse random 2 = 1 [ask turtle turtlenum [set native blue]][ask turtle turtlenum [set native green]]]
end

to magentanative [turtlenum]
let redcount count (turtles-on [neighbors] of turtle turtlenum) with [color = red]
let bluecount count (turtles-on [neighbors] of turtle turtlenum) with [color = blue]
if bluecount > redcount [ask turtle turtlenum [set native red]]
if redcount < bluecount [ask turtle turtlenum [set native blue]]
if redcount = bluecount [ifelse random 2 = 1 [ask turtle turtlenum [set native red]][ask turtle turtlenum [set native blue]]]
end

to next-simulation
save-results
ifelse replication < 10
[
set replication replication + 1
]
[
ifelse happy-percent < 80 ;Should go up to 100, but for the current version that's just chaos under all parameter settings
[
set happy-percent happy-percent + 10
set replication 1
]
[
ifelse not two-groups and not multid
[
set two-groups true
set happy-percent 0
set replication 1
]
[
ifelse not multid
[
set multid true
set two-groups false
set happy-percent 0
set replication 1
set rg-ratio 1
set rb-ratio 1
set bg-ratio 1
set multid-percent 10
]
[
ifelse multid-percent < 100
[
set multid-percent multid-percent + 10
set happy-percent 0
set replication 1
]
[
ifelse bg-ratio > 0
[
set bg-ratio 0
set rb-ratio 0
set multid-percent 10
set happy-percent 0
set replication 1
]
]
]
]
stop

```

```

        ]
      ]
    ]
  ]
]
setup-model
end

to save-results
  file-open "SegResult.txt"
  ;;File format: Preference, Twogroups, MultiID, rgratio, rbratio, bgratio, Overall, Red, Green, Blue, Cycle, Chaos
  file-type happy-percent
  file-type " "
  ifelse two-groups
  [file-type 1]
  [file-type 0]
  file-type " "
  ifelse multid
  [file-type multid-percent]
  [file-type 0]
  file-type " "
  file-type rg-ratio
  file-type " "
  file-type rb-ratio
  file-type " "
  file-type bg-ratio
  file-type " "
  file-type segregation
  file-type " "
  file-type redseg
  file-type " "
  file-type greenseg
  file-type " "
  file-type blueseg
  file-type " "
  file-type cycle
  file-type " "
  file-print chaos
  file-close
end

```

ICEM

```

globals [cycle foodmean foodsd wanderstat foragestat stealstat sharestat wanderturn forageturn stealturn sharetun wins losses]
turtles-own [health age rules priority prules vision target region tired stolen gifted rep]
patches-own [food]

```

;Rule structure:

;Num of Rules (repeat X times) Identifier Variable Comparison Constant (last step action)

;Therefore, rules structure:

```
;[[0 0][1 0 0 0 0 0][2 0 0 0 0 0 0 0 0]] etc.
```

;Identifier 0=self 1=nearestother 2=furthestother 3=strongest 4=weakest [If rep is in use 0=self 1=lowrep 2=medrep 3=highrep]

;Variable 0=age 1=health 2=food

;Comparison 0=lessthan 1=morethan 2=equal 3=wildcard

;Constant 0-100

;Action 0=wander 1=forage 2=steal 3=share

;Therefore the example rule above is "In all cases wander" then "If self age is lessthan 0 then wander" then "If self age is lessthan 0 and if self age is lessthan 0 then wander"

to setup-model

```

;; (for this model to work with NetLogo's new plotting features,
;; __clear-all-and-reset-ticks should be replaced with clear-all at
;; the beginning of your setup procedure and reset-ticks at the end
;; of the procedure.)
__clear-all-and-reset-ticks
set foodmean 0.06

```

```

set foodsd 0.001
set wanderstat [0 0 0 0] ; These four are the stats we gather for the sim!
set foragestat [0 0 0 0]
set stealstat [0 0 0 0]
set sharestat [0 0 0 0]
set wanderturn [0 0 0 0]
set forageturn [0 0 0 0]
set stealturn [0 0 0 0]
set sharetun [0 0 0 0]
set wins [0 0 0 0]
set losses [0 0 0 0]
setup-borders
populate-world
end

```

```

to setup-borders
let x 0
while [x <= 64]
[
ask patch 0 x [set pcolor blue]
ask patch 31 x [set pcolor blue]
ask patch 62 x [set pcolor blue]
ask patch x 0 [set pcolor blue]
ask patch x 31 [set pcolor blue]
ask patch x 62 [set pcolor blue]
set x x + 1
]
end

```

```

to populate-world
populate-society 1 1 300 9 1
populate-society 1 32 300 9 2
populate-society 32 1 300 9 3
populate-society 32 32 300 9 4
end

```

```

to populate-society [ecks why tnum fnum reg]
let x 0
let y 0
let remaining 900
while [x <= 29]
[
while [y <= 29]
[
if random(100) <= (tnum / remaining) * 100
[
set tnum tnum - 1
spawn-turtle (ecks + x) (why + y) reg
]
if random(100) <= (fnum / remaining) * 100
[
set fnum fnum - 1
spawn-food (ecks + x) (why + y)
]
set remaining remaining - 1
set y y + 1
]
set x x + 1
set y 0
]
end

```

```

;Identifier 0=self 1=nearestother 2=furthestother 3=strongest 4=weakest
;Variable 0=age 1=health 2=food (3= reputation if in use)
;Comparison 0=lessthan 1=morethan 2=equal 3=wildcard
;Constant 0-100
;Action 0=wander 1=forage 2=steal 3=share

```

```

to spawn-turtle [ecks why reg]
let rulenum 5 + random(3)

```

```

let rulelist []
let newrule []
let clauses 0
while [rulenum > 0]
[
  set newrule []
  set newrule lput random(3) newrule ;Provides number of clauses 0-2
  set clauses item 0 newrule
  while [clauses > 0]
  [
    ifelse reputation > 0 [set newrule lput random (4) newrule][set newrule lput random (5) newrule] ;ID (need one less if reputation is on!)
    set newrule lput random (3) newrule ;Var
    set newrule lput random (4) newrule ;Comparison
    set newrule lput random (101) newrule ;Constant
    set clauses clauses - 1
  ]
  set newrule lput random (4) newrule ;Action
  set rulelist fput newrule rulelist
  set rulenum rulenum - 1
]
ask patch ecks why
[
  sprout 1
  [
    set color [160 160 160]
    set shape "circle"
    set age 0
    set health 100
    set rules rulelist
    set vision 3
    set region reg
    set rep 50 ;Neutral, 0 is good 100 is bad
    prioritise
  ]
]
end

```

;Comparison 0=lessthan 1=morethan 2=equal 3=wildcard

to prioritise ;Called from agent level, both here during setup and later after mating

```

let ruleno length rules
let clauseloop 0
let thisrule []
let thispriority 0
let comparison 0
set priority rules
while [ruleno > 0]
[
  set ruleno ruleno - 1 ;List index to one under their length, remember
  set thisrule item ruleno rules
  set thispriority 0 ;Will remain priority 0 if there are no rules
  set clauseloop item 0 thisrule
  while [clauseloop > 0]
  [
    set clauseloop clauseloop - 1
    set comparison item (3 + (clauseloop * 4)) thisrule
    if comparison = 0 [set thispriority thispriority + 2] ;Small change, lessthan/morethan creates a requirement
    if comparison = 1 [set thispriority thispriority + 2] ;Small change, lessthan/morethan creates a requirement
    if comparison = 2 [set thispriority thispriority + 5] ;Big change, equals is a very specific operator
    if comparison = 3 [set thispriority thispriority] ;No change, clause may as well not exist since wildcard is always true
  ]
  set priority replace-item ruleno priority thispriority
]
set prules []
let ploop 20 ;This is overkill I think, the maximum priority is, in theory, ten - two clauses both using the equals operator
let rloop 0
while [ploop >= 0]
[
  set rloop 0
  while [rloop < length priority]

```



```

    [
      if item rloop priority = ploop
      [
        set prules lput item rloop rules prules
      ]
      set rloop rloop + 1
    ]
    set ploop ploop - 1
  ]
end

to spawn-food [ecks why]
ask patch ecks why
[
  set pcolor green
  set food 20
]
end

to start-simulation
create-food
turtles-mate
turtles-act
if society-competition > 0
[
  if remainder cycle society-competition = 0
  [
    societies-compete
  ]
]
turtles-age
turtles-die
stats
set cycle cycle + 1
if cycle = 20000
[
  ifelse save-results
  [
    save-data
    ifelse replications > 0
    [
      set replications replications - 1
      setup-model ;Essentially makes it perform another run
    ]
    [
      stop
      ;This is code for performing a parameter sweep
      ifelse society-competition = 0
      [
        set society-competition 1
        set replications 10
        setup-model
      ]
      [
        ifelse reputation < 4
        [
          set reputation reputation + 1
          set society-competition 0
          set replications 10
          setup-model
        ]
        [
          stop
        ]
      ]
    ]
  ]
]
let t1 item 0 wanderstat + item 0 foragestat + item 0 stealstat + item 0 sharestat
print foragestat

```

```

print t1
print "REGION 1"
type "Wander: " print item 0 wanderstat / t1
type "Forage: " print item 0 foragestat / t1
type "Steal: " print item 0 stealstat / t1
type "Share: " print item 0 sharestat / t1
let t2 item 1 wanderstat + item 1 foragestat + item 1 stealstat + item 1 sharestat
print "REGION 2"
type "Wander: " print item 1 wanderstat / t2
type "Forage: " print item 1 foragestat / t2
type "Steal: " print item 1 stealstat / t2
type "Share: " print item 1 sharestat / t2
let t3 item 2 wanderstat + item 2 foragestat + item 2 stealstat + item 2 sharestat
print "REGION 3"
type "Wander: " print item 2 wanderstat / t3
type "Forage: " print item 2 foragestat / t3
type "Steal: " print item 2 stealstat / t3
type "Share: " print item 2 sharestat / t3
let t4 item 3 wanderstat + item 3 foragestat + item 3 stealstat + item 3 sharestat
print "REGION 4"
type "Wander: " print item 3 wanderstat / t4
type "Forage: " print item 3 foragestat / t4
type "Steal: " print item 3 stealstat / t4
type "Share: " print item 3 sharestat / t4
stop
]
]
end

```

```

to create-food
region-create-food 1 1 random-normal foodmean foodsd
region-create-food 1 32 random-normal foodmean foodsd
region-create-food 32 1 random-normal foodmean foodsd
region-create-food 32 32 random-normal foodmean foodsd
end

```

```

to region-create-food [ecks why quantity]
set quantity quantity * 900
set quantity int quantity
let x 0
let y 0
let remaining 900
while [x <= 29]
[
while [y <= 29]
[
if random(100) <= (quantity / remaining) * 100
[
set quantity quantity - 1
more-food (ecks + x) (why + y)
]
set remaining remaining - 1
set y y + 1
]
set x x + 1
set y 0
]
end

```

```

to more-food [ecks why]
ask patch ecks why
[
set food food + 20
set pcolor green
]
end

```

```

to turtles-mate
ask turtles
[

```

```

if age > 18 and health > 50
[
  if suitable-mate
  [
    mate one-of spotted-turtles with [age > 18 and health > 50]
  ]
]
]
end

to mate [partner] ;Called at turtle level, partner is reporting a turtle (not a who number)
let options legal-patches partner
if is-agentset? options
[
  if count options > 0
  [
    let r1 rules
    let r2 []
    let newrule []
    let clauses 0
    let p1rule []
    let p2rule []
    let p1clause []
    let p2clause []
    let clauseloop 0
    let cmin 0
    let cmax 0
    let mutantclause []
    let cloop 0
    ask partner [set r2 rules]
    if length r2 > length r1
    [
      let r3 r1
      set r1 r2
      set r2 r3
    ]
    let i 0
    let i2 0
    let parentrules 0
    while [i < length r2]
    [
      ;Iterating over each of the rulesets
      set newrule []
      ;So now the newrule just has the number of clauses specified...due to the mutation chance this might be more clauses than the parents
      have?
      set p1rule item i rules
      ifelse length r2 < i
      [
        set p2rule p1rule
      ]
      [
        ask partner [set p2rule item i rules]
      ]
      ;Setup number of clauses
      ifelse random(100) <= mutation-rate
      [
        set clauses random(2) ;Bear in mind this may mean that the rule has more clauses than either parents rule!
      ]
      [
        ifelse random(2) = 1
        [
          set clauses item 0 p1rule
        ]
        [
          set clauses item 0 p2rule
        ]
      ]
      set newrule lput clauses newrule
      set clauseloop 1
      while [clauseloop <= clauses]

```

```

[
  set cmin 1 + ((clauseloop - 1) * 4) ;Clause 1 starts at item 1, Clause 2 starts at item 5
  set cmax cmin + 4 ;Clauses are four items long, so X, X+1, X+2 and X+3, sublist treats the second number as exclusive so it has to
be one higher
  ifelse item 0 p1rule >= clauseloop
  [
    ifelse item 0 p2rule >= clauseloop
    [
      ;Combine clauses from two parents
      set p1clause sublist p1rule cmin cmax
      set p2clause sublist p2rule cmin cmax
    ]
    [
      ;Only p1 has an appropriate clause
      set p1clause sublist p1rule cmin cmax
      set p2clause sublist p1rule cmin cmax
    ]
  ]
  [
    ifelse item 0 p2rule >= clauseloop
    [
      ;Only p2 has an appropriate clause
      set p1clause sublist p2rule cmin cmax
      set p2clause sublist p2rule cmin cmax
    ]
    [
      ;Neither parent has an appropriate clause - MUTANT CLAUSE!
      set mutantclause []
      ifelse reputation > 0 [ set mutantclause lput random(4) mutantclause][ set mutantclause lput random(5) mutantclause]
      set mutantclause lput random(3) mutantclause
      set mutantclause lput random(4) mutantclause
      set mutantclause lput random(101) mutantclause
      set p1clause mutantclause
      set p2clause mutantclause
    ]
  ]
]
;Now p1clause and p2clause are both complete clauses containing info to be combined to create the new clause
set cloop 0
while [cloop < length p1clause]
[
  ifelse random(100) <= mutation-rate
  [
    let mutatemax 0
    if cloop = 0 [ifelse reputation > 0 [set mutatemax 4][set mutatemax 5]]
    if cloop = 1 [set mutatemax 3]
    if cloop = 2 [set mutatemax 4]
    if cloop = 3 [set mutatemax 101]
    set newrule lput random(mutatemax) newrule
  ]
  [
    ifelse random(2) = 1
    [
      set newrule lput item cloop p1clause newrule
    ]
    [
      set newrule lput item cloop p2clause newrule
    ]
  ]
  set cloop cloop + 1
]
set clauseloop clauseloop + 1
]
;Getting there! At this point "newrule" contains the number of rules and all of the clauses. It just needs an action to finish it off
ifelse random(100) < mutation-rate
[
  set newrule lput random(4) newrule
]
[
  ifelse random(2) = 1
  [

```

```

        set newrule lput last p1rule newrule
      ]
    [
      set newrule lput last p2rule newrule
    ]
  ]
  ;Finally newrule has the new rule within it. Reusing the r1 variable to contain the new ruleset (we'll never look at this item of r1 again
so we can scrap it)
  set r1 replace-item i r1 newrule
  set i i + 1
]
;At the end of all of this r1 should contain the new ruleset. Newrule only contains each line of the rule one at a time
let newhealth round (0.25 * health)
set health health - newhealth
ask partner
[
  let p2gift round (0.25 * health)
  set health health - p2gift
  set newhealth newhealth + p2gift
]
let thisreg region
ask one-of options
[
  sprout 1
  [
    set color [160 160 160]
    set shape "circle"
    set age 0
    set health newhealth
    set rules r1
    set vision 3
    set rep 50 ;Neutral, 0 is good 100 is bad
    set region thisreg
    prioritise
  ]
]
set tired 1
]
stop

]
end

to-report legal-patches [partner]
  let options []
  set options (patch-set [neighbors4] of patch-here options)
  ask partner [set options (patch-set [neighbors4] of patch-here options)]
  if is-agentset? options
  [
    set options options with [count turtles-here = 0 and not (pcolor = blue)]
  ]
  report options
end

to-report suitable-mate ;Run from the turtle level!
  ifelse is-agentset? spotted-turtles
  [
    ifelse count spotted-turtles with [age > 18 and health > 50] > 0
    [report true]
    [report false]
  ]
  [
    report false
  ]
]
end

to-report spotted-turtles ;Run from the agent level!
  let spotted []
  let ecks xcor
  let why ycor

```

```

let look 1
while [look <= vision]
[
  ifelse [pcolor] of patch (ecks + look) why = blue
  [
    set look vision
  ]
  [
    if count turtles-on patch (ecks + look) why > 0
    [
      set spotted (turtle-set turtles-on patch (ecks + look) why spotted)
    ]
  ]
  set look look + 1
]
set look -1
while [look >= (0 - vision)]
[
  ifelse [pcolor] of patch (ecks + look) why = blue
  [
    set look 0 - vision
  ]
  [
    if count turtles-on patch (ecks + look) why > 0
    [
      set spotted (turtle-set turtles-on patch (ecks + look) why spotted)
    ]
  ]
  set look look - 1
]
set look 1
while [look <= vision]
[
  ifelse [pcolor] of patch ecks (why + look) = blue
  [
    set look vision
  ]
  [
    if count turtles-on patch ecks (why + look) > 0
    [
      set spotted (turtle-set turtles-on patch ecks (why + look) spotted)
    ]
  ]
  set look look + 1
]
set look -1
while [look >= (0 - vision)]
[
  ifelse [pcolor] of patch ecks (why + look) = blue
  [
    set look 0 - vision
  ]
  [
    if count turtles-on patch ecks (why + look) > 0
    [
      set spotted (turtle-set turtles-on patch ecks (why + look) spotted)
    ]
  ]
  set look look - 1
]
report spotted
end

to turtles-act
ask turtles
[
  ifelse tired = 1
  [
    set tired 0 ;If a turtle mated this round it doesn't get to take an action
  ]
]

```

```

[
  let checking 0
  while [checking < length prules]
  [
    if cando item checking prules
    [
      dodo last item checking prules
      set checking length prules + 999
    ]
    set checking checking + 1
  ]
  if checking = length prules
  [dodo 1] ;Default to foraging if all rules say nothing - as in original simulation
]
]
end

```

to-report cando [rule] ;Is called from turtle level

```

let clauses item 0 rule
let clause []
if clauses = 0 [set target self]
while [clauses > 0]
[
  set clause sublist rule (1 + ((clauses - 1) * 4)) (5 + ((clauses - 1) * 4)) ;Clause 1 is items 1-4, sublist is exclusive so equation needs to
report 1-5. C2 is 5-9
  if not trueclause clause [report false]
  set clauses clauses - 1
]
report true
end

```

;Reminder on rule format...

;Identifier 0=self 1=nearestother 2=furthestother 3=strongest 4=weakest

;Variable 0=age 1=health

;Comparison 0=lessthan 1=morethan 2=equal

;Constant 0-100

to-report trueclause [clause] ;Still on the turtle level...

```

let identifier item 0 clause
let variable item 1 clause
let comparison item 2 clause
let constant item 3 clause
if identifier = 0 [set target self]
ifelse identifier > 1 and is-agentset? spotted-turtles
[
  ifelse reputation > 0
  [
    ;If reputation is on identify agents by reputation
    if identifier = 1
    [
      ifelse count spotted-turtles with [rep < 50] > 0
      [
        set target one-of spotted-turtles with [rep < 50]
      ]
      [
        report false
      ]
    ]
  ]
  if identifier = 2
  [
    ifelse count spotted-turtles with [rep = 50] > 0
    [
      set target one-of spotted-turtles with [rep = 50]
    ]
    [
      report false
    ]
  ]
  if identifier = 3
  [

```

```

    ifelse count spotted-turtles with [rep > 50] > 0
    [
      set target one-of spotted-turtles with [rep > 50]
    ]
    [
      report false
    ]
  ]
]
[
  ;If reputation is off identify agents by distance or health
  if identifier = 1 [set target min-one-of spotted-turtles [distance myself]]
  if identifier = 2 [set target max-one-of spotted-turtles [distance myself]]
  if identifier = 3 [set target max-one-of spotted-turtles [health]]
  if identifier = 4 [set target min-one-of spotted-turtles [health]]
]
]
[
  report false
]
let leftside 0
if variable = 0
  [set leftside [age] of target]
if variable = 1
  [set leftside [health] of target]
if variable = 2 ;This tests how many vision squares have food
[
  set leftside count spotted-patches with [food > 0]
  if [food] of patch-here > 0 [set leftside leftside + 1]
]
let rightside constant
if comparison = 0
[
  ifelse leftside < rightside
  [report true][report false]
]
if comparison = 1
[
  ifelse leftside > rightside
  [report true][report false]
]
if comparison = 2
[
  ifelse leftside = rightside
  [report true][report false]
]
if comparison = 3
[
  report true ;The wildcard comparison is always evaluated as true
]
if comparison > 3 [print who print clause print "WTF!!!"]
end

to dodo [action] ;Is called from the turtle level
if action = 0 [wander]
if action = 1 [forage]
if action = 2 [steal]
if action = 3 [give]
end

to wander ;Agent moves. Action fails if no valid patch is present
set wanderturn replace-item (region - 1) wanderturn (item (region - 1) wanderturn + 1)
let options spotted-patches
if is-agentset? options
[
  set options options with [count turtles-here = 0]
  if count options > 0
  [
    move-to one-of options
  ]
]

```



```

]
eat
end

```

to-report spotted-patches ;Run from the agent level!

```

let spotted []
let ecks xcor
let why ycor
let look 1
while [look <= vision]
[
  ifelse [pcolor] of patch (ecks + look) why = blue
  [
    set look vision
  ]
  [
    set spotted (patch-set patch (ecks + look) why spotted)
  ]
  set look look + 1
]
set look -1
while [look >= (0 - vision)]
[
  ifelse [pcolor] of patch (ecks + look) why = blue
  [
    set look 0 - vision
  ]
  [
    set spotted (patch-set patch (ecks + look) why spotted)
  ]
  set look look - 1
]
set look 1
while [look <= vision]
[
  ifelse [pcolor] of patch ecks (why + look) = blue
  [
    set look vision
  ]
  [
    set spotted (patch-set patch ecks (why + look) spotted)
  ]
  set look look + 1
]
set look -1
while [look >= (0 - vision)]
[
  ifelse [pcolor] of patch ecks (why + look) = blue
  [
    set look 0 - vision
  ]
  [
    set spotted (patch-set patch ecks (why + look) spotted)
  ]
  set look look - 1
]
report spotted
end

```

to forage ;Agent goes for the food. Action fails if there are no free food patches, but may still eat

```

set foragerturn replace-item (region - 1) foragerturn (item (region - 1) foragerturn + 1)
let options spotted-patches
if is-agentset? options
[
  set options options with [count turtles-here = 0]
  if count options > 0
  [
    set options options with [food > 0]
    if count options > 0
    [

```

```

        move-to one-of options
    ]
]
]
eat
end

to eat ;Agent noms down available food
ifelse [food] of patch-here > 20
[
    set health health + 20
    ask patch-here [set food food - 20]
]
[
    set health health + [food] of patch-here
    ask patch-here [set food 0 set pcolor black]
]
end

```

```

to steal ;Agent steals. If no target is available action is wasted.
set stealturn replace-item (region - 1) stealturn (item (region - 1) stealturn + 1)
if reputation = 1
[
    set stolen stolen + 1
    if stolen > gifted [set rep 100]
    if stolen = gifted [set rep 50]
]
if reputation = 3
[
    set rep 100
]
if target = self
[
    if is-agentset? spotted-turtles
    [
        set target one-of spotted-turtles
    ]
]
ifelse target = self ;i.e. if target was self a moment ago and there's nobody new
[ ];No action!
[
    ifelse [health] of target > 25
    [
        set health health + 25
        ask target [set health health - 25]
    ]
    [
        set health health + [health] of target
        ask target [die]
    ]
]
end

```

```

to give ;Agent gives. If no target available action is wasted
set shareturn replace-item (region - 1) shareturn (item (region - 1) shareturn + 1)
if reputation = 1
[
    set gifted gifted + 1
    if stolen < gifted [set rep 0]
    if stolen = gifted [set rep 50]
]
if reputation = 2
[
    set rep 0
]
if target = self
[
    if is-agentset? spotted-turtles
    [
        set target one-of spotted-turtles
    ]
]
end

```

```

]
]
ifelse target = self
[] ;Again no action if no target found
[
  let gift round health * 0.25
  set health health - gift
  ask target [set health health + gift]
]
end

```

to societies-compete

if ongoing-stats ;Perform graph drawing here instead so graphs match what actually happens! Deaths and births can mess this a little

```

[
  set-current-plot "Health"
  set-current-plot-pen "Region 1"
  plot sum [health] of turtles with [region = 1]
  set-current-plot-pen "Region 2"
  plot sum [health] of turtles with [region = 2]
  set-current-plot-pen "Region 3"
  plot sum [health] of turtles with [region = 3]
  set-current-plot-pen "Region 4"
  plot sum [health] of turtles with [region = 4]
]
let fitness []
set fitness lput sum [health] of turtles with [region = 1] fitness
set fitness lput sum [health] of turtles with [region = 2] fitness
set fitness lput sum [health] of turtles with [region = 3] fitness
set fitness lput sum [health] of turtles with [region = 4] fitness
let winner 1
let loser 1
let x 1
while [x <= 3]
[
  if item x fitness > item (winner - 1) fitness [set winner x + 1]
  if item x fitness < item (loser - 1) fitness [set loser x + 1]
  set x x + 1
]
set wins replace-item (winner - 1) wins (item (winner - 1) wins + 1)
set losses replace-item (loser - 1) losses (item (loser - 1) losses + 1)
ask turtles with [region = winner]
[
  set health health + 1
]
ask turtles with [region = loser]
[
  set health health - 1
  if health <= 0 [die]
]
end

```

to turtles-age

```

ask turtles
[
  set age age + 1
  if age > 100 [die]
]
end

```

to turtles-die

```

ask turtles
[
  set health health - 10
  if health <= 0 [die]
]
end

```

to stats

```

let i 0
if cycle > 15000

```

```

[
  while [i1 < 4]
  [
    set wanderstat replace-item i1 wanderstat (item i1 wanderstat + item i1 wanderturn)
    set foragestat replace-item i1 foragestat (item i1 foragestat + item i1 forageturn)
    set stealstat replace-item i1 stealstat (item i1 stealstat + item i1 stealturn)
    set sharestat replace-item i1 sharestat (item i1 sharestat + item i1 sharetturn)
    set i1 i1 + 1
  ]
]
if ongoing-stats
[
  graphs
]
set wanderturn [0 0 0 0]
set forageturn [0 0 0 0]
set stealturn [0 0 0 0]
set sharetturn [0 0 0 0]
end

```

```

to graphs
  set-current-plot "Region One"
  set-current-plot-pen "Population"
  plot count turtles with [region = 1]
  set-current-plot-pen "Wander"
  plot item 0 wanderturn
  set-current-plot-pen "Forage"
  plot item 0 forageturn
  set-current-plot-pen "Steal"
  plot item 0 stealturn
  set-current-plot-pen "Give"
  plot item 0 sharetturn
  set-current-plot "Region Two"
  set-current-plot-pen "Population"
  plot count turtles with [region = 2]
  set-current-plot-pen "Wander"
  plot item 1 wanderturn
  set-current-plot-pen "Forage"
  plot item 1 forageturn
  set-current-plot-pen "Steal"
  plot item 1 stealturn
  set-current-plot-pen "Give"
  plot item 1 sharetturn
  set-current-plot "Region Three"
  set-current-plot-pen "Population"
  plot count turtles with [region = 3]
  set-current-plot-pen "Wander"
  plot item 2 wanderturn
  set-current-plot-pen "Forage"
  plot item 2 forageturn
  set-current-plot-pen "Steal"
  plot item 2 stealturn
  set-current-plot-pen "Give"
  plot item 2 sharetturn
  set-current-plot "Region Four"
  set-current-plot-pen "Population"
  plot count turtles with [region = 4]
  set-current-plot-pen "Wander"
  plot item 3 wanderturn
  set-current-plot-pen "Forage"
  plot item 3 forageturn
  set-current-plot-pen "Steal"
  plot item 3 stealturn
  set-current-plot-pen "Give"
  plot item 3 sharetturn
end

```

```

to save-data
  file-open "EvoSimResults.txt"
  ;Format Competition Reputation R1Wander R1Forage R1Steal R1Share R1Win R1Lose R1Dom R2Wander R2Forage ... R4Win R4Lose

```

```

R4Dom
file-type society-competition
file-type " "
file-type reputation
file-type " "
let t1 item 0 wanderstat + item 0 foragestat + item 0 stealstat + item 0 sharestat
file-type (precision (item 0 wanderstat / t1 * 100) 2)
file-type " "
file-type (precision (item 0 foragestat / t1 * 100) 2)
file-type " "
file-type (precision (item 0 stealstat / t1 * 100) 2)
file-type " "
file-type (precision (item 0 sharestat / t1 * 100) 2)
file-type " "
file-type item 0 wins
file-type " "
file-type item 0 losses
file-type " "
file-type (item 0 wins - item 0 losses)
file-type " "
let t2 item 1 wanderstat + item 1 foragestat + item 1 stealstat + item 1 sharestat
file-type (precision (item 1 wanderstat / t2 * 100) 2)
file-type " "
file-type (precision (item 1 foragestat / t2 * 100) 2)
file-type " "
file-type (precision (item 1 stealstat / t2 * 100) 2)
file-type " "
file-type (precision (item 1 sharestat / t2 * 100) 2)
file-type " "
file-type item 1 wins
file-type " "
file-type item 1 losses
file-type " "
file-type (item 1 wins - item 1 losses)
file-type " "
let t3 item 2 wanderstat + item 2 foragestat + item 2 stealstat + item 2 sharestat
file-type (precision (item 2 wanderstat / t3 * 100) 2)
file-type " "
file-type (precision (item 2 foragestat / t3 * 100) 2)
file-type " "
file-type (precision (item 2 stealstat / t3 * 100) 2)
file-type " "
file-type (precision (item 2 sharestat / t3 * 100) 2)
file-type " "
file-type item 2 wins
file-type " "
file-type item 2 losses
file-type " "
file-type (item 2 wins - item 2 losses)
file-type " "
let t4 item 3 wanderstat + item 3 foragestat + item 3 stealstat + item 3 sharestat
file-type (precision (item 3 wanderstat / t4 * 100) 2)
file-type " "
file-type (precision (item 3 foragestat / t4 * 100) 2)
file-type " "
file-type (precision (item 3 stealstat / t4 * 100) 2)
file-type " "
file-type (precision (item 3 sharestat / t4 * 100) 2)
file-type " "
file-type item 3 wins
file-type " "
file-type item 3 losses
file-type " "
file-print (item 3 wins - item 3 losses)
file-close
end

```

References

- Abrams, D. Christian, J. Gordon, D. (2007). *Multidisciplinary Handbook of Social Exclusion Research*. Hoboken, NJ: Wiley.
- Abrams, D., & Hogg, M. (1990). *Social identification, Self-categorisation and Social Influence*. Chichester: John Wiley & Sons.
- Allen, C. (2004). The Dunbar Number as a Limit to Group Sizes. In *Life with Alacrity*. Retrieved from http://www.lifewithalacrity.com/2004/03/the_dunbar_num.html
- Allport, G. W. (1954). *Nature of prejudice*. Cambridge 42, Massachusetts: Addison-Wesley Publishing Company, Inc.
- Amir, Y. (1969). Contact Hypothesis in Ethnic Relations. *Psychological Bulletin*, 71 (5), 319-342.
- Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, 193, 31-35.
- Axelrod, R. (1997). The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *The Journal of Conflict Resolution*, 41(2), 203-226.
- Badcock, P. P. (2012). Evolutionary Systems Theory: A Unifying Meta-Theory of Psychological Science. *Review of General Psychology*, 16(1), 10-23.
- Binder, J., Zagefka, H., Brown, R., Funke, F., Kessler, T., Mummendy, & Leyens, J. (2009). Does Contact Reduce Prejudice or Does Prejudice Reduce Contact? A Longitudinal Test of the Contact Hypothesis Among Majority and Minority Groups in Three European Countries. *Journal of Personality and Social Psychology*, 96 (4), 843-856.
- Bond, R. (2005). Group Size and Conformity. *Group Processes and Intergroup Relations*, 8(4), 331-354.
- Bosse, T., Gerritsen, C., & Treur, J. (2011). Combining rational and biological factors in virtual agent decision making. *Applied Intelligence*, 24, 87-101.
- Baranscombe, N. R., Schmitt, M. T., & Harvey, R. D. (1999). Perceiving Pervasive Discrimination Among African Americans: Implications for Group Identification and Well-Being. *Journal of Personality and Social Psychology*, 77 (1), 135-149.
- Brewer, M. B. (2000). Reducing prejudice through cross-categorization: Effects of multiple social identities. Reducing prejudice and discrimination. In S. Oskamp (Ed.), *Reducing Prejudice and Discrimination* (pp. 165-184). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.
- Brewer, M.B. (2003). Optimal Distinctiveness, Social Identity, and the Self. In M. Leary and J. Tangney (Eds.), *Handbook of Self and Identity*. (pp 480–491). New York, London: The Guilford Press.
- Brewer, M.B. Ho, H., Lee, J., & Miller, N. (1987). Social Identity and Social Distance Among Hong Kong Schoolchildren. *Personality and Social Psychology Bulletin*, 13 (2), 156-165.
- Clark, R., Anderson, N. B., Clark, V. R., & Williams, D. R. (1999). Racism as a Stressor for African Americans. *American Psychologist*, 54 (10), 805-816.
- Clark, W. A. V., & Fossett, M. (2008). Understanding the social context of the Schelling segregation model. *PNAS*, 105 (11), 4109-4114.
- Christ, O., Hewstone, M., Tausch, N., Wagner, U., Voci, A., Huges, J., & Cairnes, E. (2010). Direct contact as a moderator of extended contact effects: Cross sectional and longitudinal impact on outgroup attitudes, behavioural intentions, and attitude certainty. *Personality and Social Psychology Bulletin*, 36(12), 1662-1674.

- Chu, D., & Griffey, D. (1985). The contact theory of racial integration: The case of sport. *Sociology of Sport Journal*, 2, 323-333.
- Corten, R., & Buskens, V. (2010). Co-evolutions of conventions and networks: An experimental study. *Social Networks*, 32, 4-15.
- Cowan, G., & Ullman, B. J. (2006). Ingroup Rejection Among Women: The Role of Personal Inadequacy. *Psychology of Women Quarterly*, 30, 399-409.
- Cracolici, M. F., Cuffaro, M., & Nijkamp, P. (2009). A spatial analysis on Italian unemployment differences. *Statistical Methods and Applications*, 18, 275-291.
- Crisp, R. J. & Turner, R. N. (2009). Can Imagined Interactions Produce Positive Perceptions? *American Psychologist*, 64 (4), 231-240
- Crisp, R. J., Ensari, N., Hewstone, M., & Miller, N. (2002) A dual route model of crossed categorization effects. In W. Stroebe & M. Hewstone (Eds.), *European Review of Social Psychology* (Vol 13, pp 35-74). Hove, UK & Philadelphia, PA: Psychology Press.
- DeGroot, M. (1974). Reaching a Consensus. *Journal of the American Statistical Association*, 69 (345), 118-121.
- Desforges, D. M., Lord, C. G., Pugh, M. A., Sia, T. L., Scarberry, N. C., & Ratcliff, C. D. (1997). Role of group representativeness in the generalization part of the contact hypothesis. *Basic and Applied Social Psychology*, 19 (2), 183-204
- Dunbar, I. R. (1993). Coevolution of neocortical size, group size and language in humans. *Behavioral and brain sciences*, 16, 681-735.
- Dunbar, R. I. M. (2008). The social brain hypothesis. *Evolutionary Anthropology*, 6(5), 178-190.
- Dur, R., & Sol, J. (2010). Social interaction, co-worker altruism, and incentives. *Games and Economic Behaviour*, 69, 293-301.
- Eiben, A. E., & Smith, J. E. (2003). *Introduction to Evolutionary Computing*. Berlin Heidelberg: Springer-Verlag.
- Eiser, J. R., Stafford, T., & Fazio, R. H. (2009). Prejudiced learning: A connectionist account. *British Journal of Psychology*, 100, 399-413.
- Evans-Pritchard, E. E. (1940). *The Neur*. London: Oxford University Press.
- Farley, B., & Clark, W. (1954). Simulation of self-organizing systems by digital computer. *Transactions of the IRE Professional Group on Information Theory*, 4 (4), 76-84.
- Finklestein, J., Lapshin, O., & Wasserman, E. (2008). Randomised Study of Different Anti-Stigma Media. *Patient Education and Counseling*, 71, 204-214.
- Fowler, H. J., Settle, E. J., & Christakis, A. N. (2010). Correlated genotypes in friendship networks. *Proceedings of the National Academy of Sciences of the United States of America*, 108(5), 1993-1997.
- Frost, M. D. (2011). Stigma and Intimacy in Same Sex Relationships: A Narrative Approach. *Journal of Family Psychology*, 25(1), 1-10.
- Fry, P. D. (2009). Anthropological Insights for Creating Non-Warring Social Systems. *Journal of Aggression, Conflict and Peace Research*, 1(2), 4-15.

- Furtado, V., Melo, A., Coelho, L. A., Menezes, R., & Perrone, R. (2009). A bio-inspired crime simulation model. *Decision Support Systems*, 48, 282-292.
- Gill, D., & Stone, R. (2010). Fairness and desert in tournaments. *Games and Economic Behaviour*, 69, 346-364.
- Gong, T. (2010). Exploring the Roles of Horizontal, Vertical, and Oblique Transmissions in Language Evolution. *Adaptive Behavior*, 18, 356-376.
- Gotts, M. N., & Polhill, G. (2010). Size Matters: Large-Scale Replications of Experiments with FEARLUS. *Advances in Complex Systems*, 13(4), 453-467.
- Hamilton, W. D. (1964). The Genetical Evolution of Social Behavior. *Journal of Theoretical Biology*, 7(1), 1-16.
- Harrell, S. P. (2000). A Multidimensional Conceptualization of Racism-Related Stress: Implications for the Well-Being of People of Color. *American Journal of Orthopsychiatry*, 70 (1), 42-57.
- Hegselmann, R. & Kraus, U. (2002). Opinion Dynamics and Bounded Confidence Models, Analysis, and Simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 1-33
<<http://jasss.soc.surrey.ac.uk/5/3/2.html>>
- Henry, D. A., Pralat, P., & Zhang, C. (2011). Emergence of segregation in evolving social networks. *Proceedings of the National Academy of Sciences*, 108(21), 8605-8610.
- Hewstone, M., Islam, M. R., & Judd, C. M. (1993). Models of Crossed Categorization and Ingroup Relations. *Journal of Personality and Social Psychology*, 64 (5), 779-793.
- Hill, R. A., & Dunbar, R. I. M. (2002). Social network size in humans. *Human Nature*, 14(1), 53-72.
- Hodson, G. (2008). Interracial prison contact: The pros for (socially dominant) cons. *British Journal of Social Psychology*, 47, 325-351.
- Hogg, M.A. & Abrams, D. (1988). *Social Identification: A Social Psychology of Intergroup Relations and Group Processes*, London: Routledge.
- Huang, G., Saratchandran, P., & Sundararajan, N. (2005). A Generalized Growing and Pruning RBF (GGAP-RBF) Neural Network for Function Approximation. *IEEE Transactions on Neural Networks*, 16(1), 57-67.
- Iannaccone, L. R., & Makowsky, M. D. (2007). Accidental Atheists? Agent-Based Explanations for the Persistence of Religious Regionalism. *Journal for the Scientific Study of Religion*, 46(1), 1-16
- Jackman, M. R., & Crane, M. (1986). "Some of my best friends are black...": Interracial friendships and white racial attitudes. *Public Opinion Quarterly*, 50, 459-486.
- Jetten, J., Branscombe, N. R., Schmitt, M., & Spears, R. (2001). Rebels with a Cause: Group Identification as a Response to Perceived Discrimination From the Mainstream. *Personality and Social Psychology Bulletin*, 29 (9), 1204-1213.
- Jetten, J., Branscombe, R. N., Spears, R., & McKimme, M. B. (2003). Predicting the Paths of Peripherals: The Interaction of Identification and Future Possibilities. *Personality and Social Psychology Bulletin*, 29(1), 130-140.
- Jussim, L. (1986). *Self-fulfilling prophecies: A theoretical and intergrative review*, 93(4), 429-445.
- Jussim, L., Palumbo, P., Chatman, C., Madon, S., & Smith, A. (2000). Stigma and self-fulfilling prophecies. In T. Heatherton, R. Kleck, M. Hebl & J. Hull (Eds.), *The Social Psychology of Stigma* (374-418). New York, London: The Guilford Press.
- Kalik, S. M., & Hamilton, T. E. (1986). The Matching Hypothesis Re-examined. *Journal of Personality and*

Social Psychology, 51(4), 673-682

Kameda, T., & Nakanishi, D. (2002). Cost-Benefit analysis of social/cultural learning in a nonstationary uncertain environment. *Evolution and Human Behavior*, 23, 374-393.

Kim, S., Taber, C. S., Lodge, M. (2009). A Computational Model of the Citizen as Motivated Reasoner: Modelling Dynamics of the 2000 Presidential Election. *Political Behaviour*, 32, 1-28.

Lammers, J. H., Warburton, K., & Cribb, B. W. (2010). Individual competition and the emergence of density-dependant exposure to predation: how behaviours influence population processes. *Adaptive Behaviour*, 18(2), 171-183.

Lakhani, N. (2008). *Does multiple-identification protect self-esteem in low status groups? Revisiting the rejection-identification model.* (Unpublished doctoral dissertation). University of Birmingham.

Lane, V. R., & Scott, S. G. (2007). The neural network model of organization identification. *Organizational Behaviour and Human Decision Processes*, 104, 175-192.

Li, P. (2008). Lexical Organisation and Competition in First and Second Languages: Computational and Neural Mechanisms. *Cognitive Science*, 33, 629-664.

Lotka, A. J. (1920), Analytical Note on Certain Rhythmic Relations in Organic Systems. The Proceedings of the Natural Academic of Science U.S., 6, 410-415.

Major, B., & Eccleston, C. P. (2005). Stigma and Social Exclusion. In D. Abrams, M. Hogg & J. Marques (Eds.), *The social psychology of inclusion and exclusion* 63-88. New York: Psychology Press.

Major, B. & O'Brien, L. T. (2005). The Social Psychology of Stigma. *Annual Psychology Review*, 56, 393-421

McCarty, C., Killworth, D. P., Bernard, R. H., Johnsen, C. E., & Shelly, A. G. (2001). Comparing Two Methods for Estimating Network Size. *Human Organization*, 60(1), 28-39.

McCoy, S. K., & Major, B. (2003). Group Identification Moderates Emotional Responses to Perceived Prejudice. *Personality and Social Psychology Bulletin*, 29 (8), 1005-1017.

Merilaita, S., & Tullberg, S. B. (2005). Constrained Camouflage Facilitates the Evolution of Conspicuous Warning Coloration. *Evolution*, 59(1), 38-45.

Muller, G., & Vercouter, L. (2010). L.I.A.R.: Achieving Social Control in Open and Decentralised Multi-Agent Systems. *Applied Artificial Intelligence*, 24(8), 723-768.

Nahum, R. J., Harding, N. B., & Kerr, B. (2011). Evolution of Restraint in a Structured Rock-Paper-Scissors Community. *Proceedings of the National Academy of Sciences of the United States of America*, 108(S2), 10831-10838.

Nail, R. P., & Helton, B. W. (1999). On the Distinction between behavioral contagion, conversion conformity, and compliance conformity. *North American Journal of Psychology*, 1(1), 87-94.

Naor, M., & Milgram, R. M. (1980). Two preservice strategies for preparing regular class teachers for mainstreaming. *Exceptional Children*, 47 (2), 126-129.

Oakes, P. J., Turner, J. C. and Haslam, S. A. (1991), Perceiving people as group members: The role of fit in the salience of social categorizations. *British Journal of Social Psychology*, 30 (2), 125-144.

Oortwijn, M. B., Boekaerts, M., Vedder, P., & Fortuin, J. (2008). The impact of a cooperative learning experience on pupils' popularity, noncooperativeness, and interethnic bias in multiethnic elementary schools. *Educational Psychology*, 28 (2), 211-221.

- Ord, T. (2006). Further Results on the Social Iterated Prisoners Dilemma. In *A Mirror Clear*. Retrieved from <http://www.amirrorclear.net/academic/ideas/dilemma/index.html>
- Parker, J. H. (1968). The interaction of negroes and whites in an interactive church setting. *Social Forces*, 46 (3), 359-366.
- Pettigrew, T. F. (2008). Still a long way to go: American Black-White relations today. In G. Adams, M. Biernat, M. Branscombe, R. Nyla, C. Crandall & L. Wrightsman (Eds.), *Commemorating Brown: The social psychology of racism and discrimination* (pp. 45-61). Washington, DC: American Psychological Association.
- Pettigrew, T. F., & Tropp, L. R. (2006). A Meta-Analytic Test of Intergroup Contact Theory. *Journal of Personality and Social Psychology*, 90 (5), 751-783.
- Pool, S. I., & Kochen, M. (1978). Contacts and Influence. *Social Networks*, 1, 5-51.
- Postmes, T., & Jetten, J. (2006). *Individuality and the Group: Advances in Social Identity*. London: Sage.
- Radax, W., & Rengs, B. (2010). Prospects and Pitfalls of Statistical Testing: Insights from Replicating the Demographic Prisoner's Dilemma. *Journal of Artificial Societies and Social Simulation*, 13(4), 1, <<http://jasss.soc.surrey.ac.uk/13/4/1.html>>.
- Rizzuto, T. E., LeDoux, J., & Hatala, J. P. (2009). It's not just what you know, it's who you know: Testing a model of relative importance of social networks to academic performance. *Social Psychology in Education*, 12, 175-189.
- Roberts, S. G. B., Wilson, R., Fedurek, P., & Dunbar, R. I. M. (2007). Individual differences and personal social network size and structure. *Personality and Individual Differences*, 44, 954-964.
- Sani, F., & Pugliese, C. A. (2008). In the Name of Mussolini: Explaining the Schism in an Italian Right-Wing Political Party. *Group Dynamics: Theory, Research and Practice*, 12 (3), 243-253.
- Sato, K. (2010). Microsimulation of the Temporary Housing Situation Following an Urban Disaster: Case Study of an Anticipated Tokyo Metropolitan Earthquake. *Social Science Computer Review*, 29(1), 103-126.
- Schelling, T. S. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, 1, 143-186
- Schwartz, L. K., & Simmons, J. P. (2001). Contact quality and attitudes towards the elderly. *Educational Gerontology*, 27(2), 127-137.
- Segbroeck, V. S., Jong, d. S., Nowe, A., Santos, C. F., & Lenaerts, T. (2010). Learning to coordinate in complex networks. *Adaptive Behavior*, 18(5), 416-427.
- Sherif, M. (1937). An Experimental Approach to the Study of Attitudes. *Sociometry*, 1(1), 90-98.
- Shin, J., & Kudhir, J. (2010). A customer management dilemma: When is it profitable to reward ones own customers? *Marketing Science*, 29(4), 671-689.
- Sidanius, J., & Pratto, F. (1999). *Social dominance: An intergroup theory of social hierarchy and oppression*. NY: Cambridge University Press.
- Slay, S. H., & Smith, A. D. (2011). Professional Identity Construction: Using narrative to understand the negotiation of professional and stigmatized cultural identities. *Human Relations*, 64(1), 85-107.
- Smith, E. R., & Conrey, F. R. (2007). Agent-Based Modelling: A New Approach for Theory Building in Social Psychology. *Personal Social Psychology Review*, 11, 87-104
- Snijders, T. A. B., Bunt, G. G., & Steglich, C. E. G. (2010). Introduction to stochastic actor based models for

network dynamics. *Social Networks*, 32, 44-60.

Sober, E., & Wilson, S. D. (1990). *Unto Others: The evolution and psychology of unselfish behavior*. Cambridge, Massachusetts: Harvard University Press.

Spronck, P., & Berendsen, B. (2009). Evolutionary Ethics in Agent Societies. *International Journal of Social Robotics*, 1, 223-232.

Steele, C. M. (1997). A Threat in the Air. How Stereotypes Shape Intellectual Identity and Performance. *American Psychologist*, 52 (6), 613-629.

Stiller, J., & Dunbar, R. I. M. (2007). Perspective-taking and memory capacity predict social network size. *Social Networks*, 29, 93-104.

Szalai, F. & Szamado, S. (2009). Honest and cheating strategies in a simple model of aggressive communication. *Animal Behaviour*, 78, 949-959.

Tajfel, H. & Turner, J. C. (1979). An Integrative Theory of Intergroup Conflict. In W. G. Austin & S. Worchel (Eds.), *The Social Psychology of Intergroup Relations*. Monterey, CA: Brooks-Cole.

Trivers, L. R. (1971). The Evolution of Reciprocal Altruism. *The Quarterly Review of Biology*, 46(1), 35-57.

Turner, J. C. (1987). *Rediscovering the social group: A self-categorization theory*. Oxford: Blackwell.

Wagner, U., Van Dick, R., Pettigrew, F. T., & Christ, O. (2003). Ethnic Prejudice in East and West Germany: The Explanatory Power of Intergroup Contact. *Group Processes and Intergroup Relations*, 6(1), 22-36.

Warneken, F., & Tomasello, M. (2009). Varieties of Altruism in Children and Chimpanzees. *Trends in Cognitive Science*, 13(9), 397-402.

Wasserman, S., & Galaskiewicz, J. (Eds.). (1994). *Advances in social network analysis: Research in the social and behavioural sciences*. London: Sage Publications.

White, T., McQuaker, S., & Salehi-Abari, A. (2008). On the importance of relational concept knowledge in referral networks. *Artificial Intelligence Review*, 29, 287-303.

Wexler, M. N. (1995). Expanding the groupthink explanation to the study of contemporary cults. *Cultic Studies Journal*, 12 (1), 49-71.

Wilson, S. D., & Sober, E. (1994). Reintroducing group selection to the human behavioral sciences. *Behavioral and brain sciences*, 17, 585-654.

Williams, K. D., & Sommer, K. L. (1997). Social ostracism by coworkers: Does rejection lead to loafing or compensation? *Personality and Social Psychology Bulletin*, 23, 693-706.

Wright, A. S. (1991). Reconceptualizing Cult Coercion and Withdrawal: A Comparative Analysis of Divorce and Apostasy. *Social Forces*, 70(1), 125-145.

Wu, S., Li, L., Wu, Z., Liang, L., Cao, H., Yan, Z., & Li, J. (2008). A Brief HIV Stigma Reduction Intervention for Service Providers in China. *AIDS Patient Care and STDs*, 22 (6), 513-520.

Yonta, M. P., & Ndoundam, R. (2009). Opinion dynamics using majority functions. *Mathematical Social Sciences*, 57, 223-244.

Zhao, J., & Costello, S. J. (2009). Computational modelling of switching behaviour in repeated gambles. *Artificial Intelligence Review*, 27, 209-222.