Modelling cultural contagion using Social Impact Theory

A thesis submitted in partial fulfilment

of the requirement for the degree of Doctor of Philosophy

Rhodri L. Morris

April 2022

Cardiff University
School of Computer Science & Informatics

To

Esther, Angharad and Arianwen

Abstract

In the context of computational social modelling, culture represents the broad set of attributes over which people may influence each other. Existing agent-based models typically have simple dyadic means by which individuals copy cultural traits from others. Building on these models, we instead embed forces from Social Impact Theory as the basis on which simulated agents influence each other. Agents are influenced based on their social status as well as their similarity to, and distance from, multiple connected contacts. These principles are gradually added to a novel agent-based model, allowing us to isolate and observe their effects.

We find that these mechanisms cause a shift away from the global polarisation of existing models, toward a diverse state where cultures can overlap and mix. Application of the new model to different network topologies offers insights into the contexts which encourage either cultural convergence or pluralism. We develop methods for the generation of generalised hierarchical structures containing 'team' subgraphs, allowing us to investigate the spread of cultural traits within hypothetical organisations. The new model is also tested against existing cultural models on a real-world dataset to determine its ability to approximate actual behavioural contagion. The results support the notion that models incorporating elements of Social Impact Theory, particularly multiple sources of influence, more closely resemble reality.

Our findings offer new insights into factors affecting the diffusion of traits and cultural overlap. These insights are potentially relevant to modelling the spread of misinformation, political polarisation, and the design of social and organisational networks with regard to their cultural repercussions.

Acknowledgements

I would like to begin by thanking my supervisors, Professor Roger Whitaker, Professor Alun Preece and Dr Liam Turner for their continued support and guidance. Roger, for his invaluable advice, enthusiasm, and encouragement; his experience and knowledge of the field meant I was always subtly nudged toward the right direction. Alun, for offering me guidance in presenting and demonstrating my work to a broad range of people on the DAIS-ITA project. Liam, for his guidance and valuable comments on my work, and seeming to know the school's systems and procedures (PGR and otherwise) inside-out. I have really enjoyed our discussions at our weekly meetings.

I must also thank Dr Gualtiero Colombo for his friendship, generosity and support, and for helping me navigate academia in general. Thank you to all the staff and PhD students of the School of Computer Science and Informatics who have offered feedback, help, and good humour - too numerous to name individually for fear of leaving anyone out.

I would like to thank my family for their love and support during what has sometimes been a difficult time for us all. To my parents, for all they have done for me and my family not just during this PhD but throughout my life; it is an understatement to say that this would not have been possible without you. To the rest of my extended family for all the kind words and good times. To my daughters, Angharad and Arianwen; I can't *really* say you helped, but you made it all worthwhile. Above all, I need to thank my wife, Esther; thank you for believing in me, for your love and patience, for being my best friend, for the sacrifices you have made and the time you have allowed me over the last few years. I owe you everything.

This work was supported by the UK Engineering and Physical Sciences Research Council.

Contents

List of Publications

The work introduced in this thesis is based in part upon the following publications:

- Morris, R., Turner, L., Whitaker, R. and Giammanco, C., 2019, July. The Impact of Peer Pressure: Extending Axelrod's Model on Cultural Polarisation. In *2019 IEEE International Conference on Cognitive Computing (ICCC)* (pp. 114-121). IEEE. [112]

- Morris, R., Turner, L., Whitaker, R. and Giammanco, C., 2020, October. Breadth Verses Depth: The Impact of Tree Structure on Cultural Influence. In: *Thomson R., Bisgin H., Dancy C., Hyder A., Hussain M. (eds) Social, Cultural, and Behavioral Modeling. SBP-BRiMS 2020.* Lecture Notes in Computer Science, vol 12268. Springer, Cham. [113]

- Morris, R., Turner, L., Whitaker, R. and Giammanco, C., 2022. The effect of hierarchical status and horizontal peer influence on organisational culture. Manuscript in preparation. [114]

- Bellamy, R., Colombo, G., Eshghi, S., De Mel, G., Giammanco, C., Morris, R., Rand, D.G., Turner, L.D., Whitaker, R.M. and Williams, G.R., 2018, April. A computational framework for modelling inter-group behaviour using psychological theory. In *Next-Generation Analyst VI* (Vol. 10653, p. 106530G). International Society for Optics and Photonics.[18]

List of Figures

List of Tables

List of Algorithms

List of Symbols

a_0	Often used to denote agent 0, etc.
A	The set of agents, of size n.
α	A trait value, where $\alpha \in T$.
α'	The highest scoring trait. i.e. $ts_{\alpha',k,i} \geq ts_{\alpha,k,i}$, $\forall \alpha \in T$
C	Clustering coefficient in Chapter 4.
d	The maximum possible distance of interaction (Chapter 6).
δ	Kronecker delta.
F	Number of Features. Note that while capitalisation typically denotes a set, for consistency with other related models we use F to refer to the size of the feature vector rather than the vector itself.
h	Height of tree (Chapter 5).
i, j	Typically used to denote nodes in a graph, during interaction with each other.
k	In the context of culture model definitions, used for feature number.
k	In the context of generating k-ary trees (Chapter 5), used for tree branching factor.
l	Current distance, where $l \leq d$ (Chapter 6).

L	Average shortest path in Chapter 4, for consistency with Watts/Strogatz 1998.
$N(i)$	Set of neighbours of node i.
$N^{+}(i)$	Set of out-neighbours from node i.
P	The current path from i, where $P \in Paths_i^l$ (Chapter 6).
$Paths_i^l$	Set of all paths of length l from node i (Chapter 6).
p	Used generally as a variable for probability, but particularly as the probability of an intra-team edge existing (Chapter 5) or the probability that an edge from a ring lattice is 'rewired' in a small-world network (Chapters 4 and 6).
q	Number of Traits. Equivalent to $\|T\|$.
σ_k^i	The value of the k^{th} feature of i.
$sim_{i,j}$	The similarity between i and j.
$strength_{i,j}$	The *Strength* of influence between i, j, where $strength_{i,j} = sim_{i,j} \times w_{i,j}$.
T	The set of possible traits, of size q.
τ	Threshold for contagion (Chapter 4).
$ts_{\alpha,k,i}$	Trait score for for trait α, feature k, node i.
$w_{i,j}$	Weight of edge between i, j.

Chapter 1

Introduction

1.1 Cultural Dissemination and Contagion

The effects of individuals and groups influencing others to adopt behaviour is a long-standing consideration across multiple fields from social, political and economic sciences to more recent areas such as artificial intelligence. The dissemination of values, practices, norms, beliefs and opinions, intentional or otherwise, has implications for the diffusion of innovations, the 'viral' spread of information (and misinformation), the growth of religions, stability of markets, and distribution of ancient artefacts. The willingness — or not — of recipients to accept these transmissions can affect the formation, identity, and fracture of groups and even nations.

The combination of these mutable attributes over which people may influence each other can be broadly thought of as *culture*. Those of the same, or similar, culture will share many of the same attributes. Conversely, differentiation of cultural identities may demarcate separate cultures, leading to polarisation and possibly conflict. Despite tendencies of people to learn and adopt behaviour from others, and thus become more alike [90, 103], individuals and groups retain differences. This has been considered a somewhat paradoxical phenomena, and has stimulated much academic interest. There are various social and psychological explanations for different contexts under which groups do *not* assimilate others, such as social differentiation and geographical isolation. In his influential 1997 work, *The dissemination of culture* [6], Robert Axelrod used a simple, abstract, agent-based model to show that merely the act of copying cultural traits from those most similar can counter-intuitively result in the existence of multiple cultures. However, for those cultures to persist in this model, they must become completely polarised from each other: an implication Axelrod notes as 'problematic' in resolving tensions of a multicultural society.

If polarisation — and the complete rejection of other cultural values — are requirements for the existence of multiple cultures, the implications for diplomacy and co-operation are stark. Yet differences between individuals and groups persist alongside attributes held in common. In reality, adjacent cultures do not necessarily converge to a monoculture, nor diverge to hold completely opposing traits. Subcultures form and exist within larger cultures, minority languages survive despite pressure from neighbours with which much is held in common. In this thesis, we investigate potential mechanisms by which cultural pluralism may continue to exist in a system where similarity drives convergence.

Polarisation and the spread of misinformation have received much attention in recent years [8, 10, 11, 20, 42, 156], and has arguably become more relevant to global events with technologies such as social media making it easier than ever before to communicate. The existence of apparent 'echo chambers' [24, 50] highlights the roles of social reinforcement and 'homophily' — the tendency of individuals to associate with those more similar to themselves [103]. Meanwhile, the growth of online social networks and the notion of 'going viral' have further increased the use of the word *contagion* as a metaphor for the spread of media, ideas, and behaviours [108]. Both the psychological phenomena behind why people would choose to amplify and transmit such messages, and the network structures which enhance or impede their diffusion provide rich veins for varied research. The application of network science to the study of social connections is important and well established [15, 30, 38, 69, 160]. The engineering of these networks, deliberate or otherwise, in the shaping of groups and societies, has been considered. The performance of organisations has been studied in regard to their structure [30, 33, 71]. Damon Centola notes that Hannah Arendt made observations on the social connections permitted in totalitarian regimes; often casual 'weak' ties that encourage isolation and discourage the close-knit groups capable of mounting a challenge [5, 36]. Conversely, network analysis has examined the conditions in which collective action is possible [101, 148].

The applications and implications of behavioural spread over a network need not be restricted to scenarios where the agents are human beings. As artificial intelligence and distributed systems progress, autonomous agents may need to learn beneficial and contextual behaviour from their similar neighbours. Examination of collective social behaviours has previously provided inspiration for heuristic methods [54, 60].

1.2 Embedding Social Impact Theory in Cultural Contagion

In this thesis we take the long-standing psychological theory of Social Impact [88] and explore the effects of embedding this to gain more insightful models for cultural contagion. Current culture models largely assume simple dyadic copying without regard to social reinforcement [6, 34], and typically always result in a completely polarised state. In seeking to refine these models, we look to apply alternative simple mechanisms which may produce results more closely aligned with real-world observations. Social theories with well defined interactions, particularly those represented as functions such as Social Impact Theory, lend themselves to computational agent-based modelling [18]. Building on the work of Robert Axelrod [6] as a useful baseline, we model the dissemination of abstract cultural traits while additionally introducing the psychological concept of social influence to understand the effects that it invokes. Specifically, we base the means of cultural transmission on Latané's Social Impact theory of social influence [88] which posited three key elements of social impact on a target being influenced: the *strength* and *immediacy* of an influence, and the *number* of influences[1]. It is these concepts that we use to form the mechanisms for copying. By using a theory of social influence as the basis of the trait-copying process, we are able to assess the effects that this extended psychological model of contagion produces as compared to those of Axelrod. We are also able to assess to what extent this model reflects real-world data and alternative contexts for interaction, including fundamental organisational hierarchies.

By embedding our model and its relationships and interactions in network graphs, we can create synthetic scenarios for analysis of circumstances that may engender cultural convergence or disassociation. Where influence is affected by strength, number, and distance, the structure of social or organisational networks can have a profound effect on the spread of traits. This has implications for the formation of culturally similar groups, and the diffusion of innovations, beliefs, and behaviours. Depending on context, the spread of behaviours could be seen as desirable or not. Similarly, there may be different circumstances where policymakers may wish to encourage either greater or lesser similarity in an organisation or society's members, or to protect certain traits. Agent-based social models such as ours allow investigations on such hypothetical scenarios. Our model in particular will highlight the role of multiple sources of cultural influence, and whether influence from more distant contacts is significant.

In taking this approach we consider several key questions aligned to the effects of embedding different components of Social Impact Theory in alternative contexts. Based on validation of Social Impact Theory in the literature, *we hypothesise that extending Axelrod's model of cultural contagion by explicitly embedding Social Impact Theory will provide additional insights for synthetic problem scenarios while also aligning with data from real-world observations.*

To explore this, we ask the following research questions (RQs):

RQ1 Will trait copying based on *number* of influences produce different macroscopic results to the dyadic copying of Axelrod's model?

[1] These aspects of Latané's theory are explained in greater detail in Section 2.2.1.

RQ2 Which local dynamics may offer explanations for global behaviours in our model that differ from those in Axelrod's?

RQ3 Can any differences between our model and Axelrod's be equated to the differences between complex and simple contagions[2]?

RQ4 What effect will network characteristics such as clustering have on the diffusion of traits using these models?

RQ5 Where some agents are given increased levels of influence due to their *status*, will the shape of the hierarchy affect their ability to influence their subordinates?

RQ6 Will the addition of influence over structural *distance* (i.e. the *immediacy* of social contact) affect the core behaviours of cultural diffusion?

RQ7 Given appropriate data, can computational culture models give an indication of how behaviour spreads through a real social network?

1.3 Thesis structure and contributions

To address the above questions in a systematic manner, we use the following chapters, which make the following contributions.

- **Chapter 2** *Background.* In this chapter we give a background to the practice of agent-based modelling for social simulation and describe both Axelrod's model of cultural dissemination and Latané's Social Impact Theory in detail. We outline key models in related areas of social influence and identify a gap in modelling of cultural diffusion using Social Impact Theory which we aim to address with our model in Chapter 3.

[2] Simple and complex contagions are described in more detail in Section 2.2.3.

- **Chapter 3** *An agent-based model of social impact based cultural dissemination.* This chapter contributes a new model of cultural diffusion incorporating Social Impact Theory. In this chapter, we focus on influence from multiple sources, and find that this alone can engender a new 'mixed' state of cultural overlap. As this behaviour is a defining characteristic of our model, we also contribute new metrics to measure this 'mixing' of overlapping cultures. These measures and the model introduced form the basis for several subsequent experiments in later chapters. This chapter supports questions RQ1 and RQ2.

- **Chapter 4** *Cultural diffusion as network contagion* contributes a network analysis and comparison of dyadic versus multiple-influence models at both local and global scale, using 'small world' networks to isolate the effects of network characteristics. Instead of the expected convergence under small-world conditions, we find that cultural mixing increases in the multiple-influence model. We also compare the local behaviours of the dyadic and multiple-influence models with the theory of simple and complex contagions, finding examples of similar dynamics but also instances where the theory does not neatly apply. This chapter supports questions RQ2, RQ3, and RQ4.

- **Chapter 5** *The effect of compound social influence in hierarchical structures* contributes an application of our model on the examination of tall versus wide hierarchical structures, fundamental in organisational structures. The model from Chapter 3 is extended to include the effects of hierarchical *status*: greater influence of those agents higher in the organisational structure. This is achieved by way of using directed and weighted edges to control the direction and strength of influence. We find that the stronger influence of higher nodes may be countered by the sheer number of subordinates, but under some conditions it may also be enhanced. Also contributed are generalised network structures for the purposes of modelling hierarchical organisations. This chapter supports question RQ5.

- **Chapter 6** *The influence of social distance.* The main contribution of this chapter is the examination of Latané's *immediacy* within our model, by means of diminished influence over greater network distance between nodes. Two mechanisms are presented for modelling this effect: an inverse-square method, and a method based on the strength of intermediaries along a network path. We repeat many of the experiments of Chapter 4 with this new behaviour included, and find counter-intuitive results concerning the widening of an agent's influences via distance. This chapter supports question RQ6.

- **Chapter 7** *An empirical test of our model on real-world data.* Moving beyond abstract scenarios, we validate models of social influence against real-world data. Our fully extended model - now including the effects of multiple influences from Chapter 3, directed and weighted edges from Chapter 5, and social distance from Chapter 6 - is tested alongside Axelrod's and another extension of his culture model. Using longitudinal data, we find that given the same starting conditions, the models give a good approximation of the data-set's later state. This chapter supports question RQ7.

- **Chapter 8** *Conclusions* in which we reflect on the contributions and results of the thesis and suggest proposals for future work.

Chapter 2

Background

The aim of this thesis largely concerns the application of theory originating from social psychology as the mechanisms within a computational simulation, specifically an agent-based culture model. We give a brief background of agent-based modelling in the social sciences, before focusing on the areas most pertinent to our work: Latané's Social Impact Theory and Axelrod's model of cultural dissemination. We examine several models and their suitability for applying Social Impact Theory to a model of cultural dissemination. We determine the mechanisms needed for doing so, and identify where existing models do not allow for simulating the relevant behaviours. We give a summary of these findings toward the end of the chapter in Table 2.1.

2.1 Computational simulations of social dynamics

The origins of quantitative approaches to social science most probably formed during the Enlightenment; inspired by Newton, philosophers such as Hobbes, Petty, and Hume sought to frame political and 'moral subjects' in more empirical and mechanistic terms [14, 39]. The growth of computing in the twentieth century precipitated its use a tool in a wide range of fields, including the social sciences. A comprehensive exploration of computational social science could fill several volumes [39, 56, 107]; instead we focus on the specific areas of networked agent-based models, specifically those involving social influence.

The study of social networks owes much to graph theory, from Euler's Konigsberg bridges to the coining of 'sociograms' by Moreno [39]. In the 1950s standard terms and concepts such as centrality and density were introduced [17], as well as models of graph structure and generation such as random graphs [146], 'preferential attachment' [39], and small-worlds [46, 106].

Meanwhile, John von Neumann and Stanislaw Ulam pioneered a theory of automata, laying the foundations for future cellular automata and agent-based models. In 1959, Oliver Benson pioneered the use of computer simulation in social sciences with a 'Simple Diplomatic Game Model', using variables such as national power and aggression to model likelihoods of diplomatic action and war [86].

The use of simulation as a tool can serve many purposes depending on context, from basic research and observation to hypothesis development and empirical testing where real data is difficult to obtain. Simulations can be run a large number of times with variations in parameters, allowing specific assumptions to be isolated for testing [18]. Joshua Epstein argues that the agent-based approach in particular is a form of 'generative social science', where if local interactions can be shown to generate 'macroscopic regularities', they may be candidate explanations for that global behaviour [57].

Of course, it is impossible to quantify and model every facet of human behaviour, and so simplifying abstractions and generalisations are always employed. One is often reminded of George Box's oft-quoted line: *"All models are wrong, but some are useful"* [26].

2.1.1 Agent based modelling

Formative contributions to the use of agent-based modelling of social interactions came from Coleman and Sakoda [21], but perhaps the most famous early agent-based simulation was Thomas Schelling's work on residential segregation dynamics [137].

Schelling's model places agents in a cellular grid. The grid is not full; some spaces remain empty, allowing agents to move to different sites under certain conditions. Each agent is assigned one of two immutable types, let us call them 'X' and 'O'. At each time-step, an agent determines whether it is 'satisfied' with its current location, if not, it moves to a random empty cell. The measure of satisfaction is how many of its neighbours are of the same type as the agent. If an agent assesses its surroundings and decides that too many of its neighbours are "not like us", it will move to a new area.

The surprising result of this model, was that even a relatively small preference for neighbours of the same 'type' results in a separating into segregated areas of 'X' and 'O', rather than a mix. The segregation (and perhaps polarisation) or neighbourhoods was not caused by explicit differentiation or movement away from the other type of agent, but simply by a slight preference for one's own; a recurring theme in models discussed in this chapter.

Another recurring theme in social agent-based modelling is the emergence of similarities and analogies between behaviours of social models and those of physics models. In Schelling's segregation, physicists saw parallels with concepts such as coarsening and the phase separation of emulsions [44]. The application of statistical physics to social models has proved a fruitful one [34]. Indeed, physicists have found many of the models discussed later in this chapter to have similarities with Ising's spin model of ferromagnetism [125]. While the focus of physicists has often been the classification of social dynamics against known physical phenomena, social models do not always lend themselves to neat categorisation of macro-scale dynamics, and local or micro-scale considerations have often been more appropriate.

A key dynamic of Schelling's model is that of emergent behaviour; individual proclivities do not necessarily result in obvious group behaviour (or as Schelling referred to them, *Micromotives and Macrobehavior* [138]). This principle has shaped the use of agent-based models since, and allies with a use of simulation to test the broader implications of specific isolated behaviours.

Common features of agents in agent-based models include being aware of its own state and that of its local environment, though depending on the model an agent's 'vision' may not extend far. Agents are autonomous, usually without any central authority, and can generally communicate or interact with other agents in some way [39]. Agents adapt to their local environment; in line with Axelrod's model, most of the models we examine feature *'adaptive rather than rational agents'* [6]. There are no costs or benefits, utility, fitness, or strategy as in some game-theoretic models; agents simply give and receive influence, reacting to their environment.

2.2 Social influence and contagion

Social influence is defined by the American Psychological Association as *'any change in an individual's thoughts, feelings, or behaviours caused by other people'* and encompasses or is related to several other concepts of social psychology such as interpersonal influence and social pressure [3]. Considerations of social environment and how people influence others date back to antiquity, but gathered pace in the wake of the French Revolution [131]. Violent uprisings of the 19$^{\text{th}}$ century became attributed to the power of *suggestion*. The Scottish journalist Charles Mackay popularised the term 'Madness of Crowds' in his book of similar name [96], which among other examples contributed a (likely exaggerated [65]) description of the Dutch tulip bubble as one of social mania. More academic treatments of crowd behaviour were given by Sighele [141], Tarde, and Le Bon [91] as the field of social psychology developed.

In the following century, the study of social influence widened into consideration of many social phenomena [131]. It is however interesting (at least for our purposes), that those early works seemed to describe it as a contagion. Indeed, both Sighele and Le Bon use the word in their works [91, 141]. The term *social contagion* is sometimes used in a specific sense (particularly in the context of the early work mentioned above) relating to the easy spread of behaviour or emotion in a crowd, and sometimes more generally to encompass broader diffusion of attributes. In this thesis we may use contagion to mean the possible transmission of any mutable attribute, but not that this transmission is necessarily trivial, immediate, or without consideration on the part of the individual. Distinctions between instances where behaviours and traits spread easily, and scenarios where they do not, are examined in Chapter 4.

Serge Moscovici formed a theory of minority influence, accounting for the fact that the norms and opinions of the majority are not always universally adopted [115]. While Moscovici's work contained a qualitative difference between how minority and majority influences achieved adoption, Bibb Latané's social impact theory [88] allowed for the possibility of minority influence through the same essential mechanisms as that of majority influence. Where influence consists of not only the *number* of influences, but also their relative individual *strength* and *immediacy*, a minority of strong influences can outweigh a local majority. Latané's theory has proven popular and influential, and has been used in various contexts [25, 78, 80, 119]. Its expression as a simple formula of strength, immediacy, and number positions social impact theory as a useful candidate for algorithmic modelling of social influence. Several models have mechanisms that appear similar to those of *strength, immediacy* or *number*, but few explicitly model all three. Social Impact Theory has rarely, if ever, been utilised as a method for examining the polarisation present in culture models such as Axelrod's [6]. In this thesis, we use Social Impact Theory as the basis for *how* the copying of cultural traits takes place on each iteration of a computational simulation. The theory is described in greater depth below.

2.2.1 Social Impact Theory

Bibb Latané's social impact theory [88] appears, on the surface, to be a rather simple description of people's social effect on each other. Inspired by physical forces, particularly light, Latané sets out three principles of social impact, in which source of influence can act upon a target:

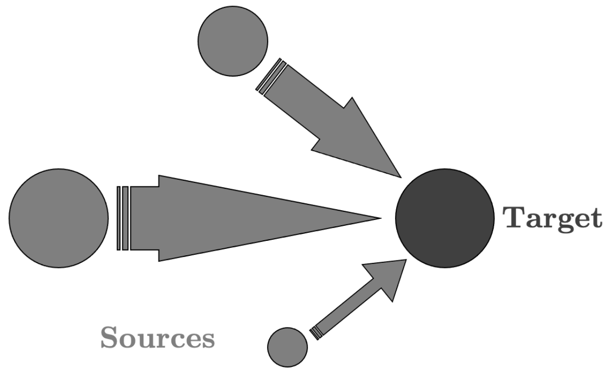

Figure 2.1: Social Impact Theory: a number of sources exerting asymmetric influence on a target, depending on their strength and distance.

- ***S**trength* - Determined by the power, importance, status, prior relationship, reputation etc. of a source of influence.
- ***I**mmediacy* - The physical or temporal proximity of the source, the closeness of contact, the absence intervening barriers or filters.
- ***N**umber* - Simply the number of sources acting upon the target.

These effects are compared by Latané to the effect of light-bulbs shining on a surface; the number of bulbs, their lumens, and their distance from the surface. He formulates these effects on impact as simply multiplicative:

$$Impact = f(SIN)$$

Thus a target's contact who has power over the target (like a boss), or is one with many things in common (perhaps a sibling), or is a source of inspiration (perhaps a celebrity), will have a high *strength* of influence. However, should that source be fairly *distant* - perhaps despite best intentions the two contacts meet less than they would like, or one lives in another country, or it may even be the case that they simply don't know each other personally (for example, the inspirational celebrity) - then the influence will be diminished. The sheer number of influences will also have an effect; one may suffer stage fright in front of an audience, or be swept along by the emotions of a crowd, or be reticent to act against the social norms of those around them.

The social force of *number* is qualified by the addition of a 'psychosocial law', that each additional source added will contribute less influence than those before. Where two people are influencing a target, the addition of a third could make a substantial difference. Where 99 people are influencing a target, the addition of a 100^{th} adds little impact. Latané expresses this:

$$Impact = sN^t$$

where N is the number of sources, s is a scaling constant, and t is a value between 0 and 1. In return, a source's impact may be diminished by the number of targets they try to influence at once:

$$Impact\ from\ source = f(1/SIN)$$

For each additional *target* that an individual seeks to influence, their impact will be a little less each time.

Latané gives supporting evidence across several scenarios, including conformity and imitation, bystander intervention, stage fright, and news interest. The theory is often cited, and has been applied to considerations of phenomena such as social loafing [80] and the influence of physical personal space [89] amongst many others [25, 78, 119].

2.2.2 Computational models of social influence and opinion formation

The key hypotheses and research questions outlined in Chapter 1 are centred on applying Social Impact Theory to culture models, as a simulation of social influence in a population. To apply Social Impact Theory computationally, it is clear that we will need to model its three main mechanisms:

- a form of *strength* of influence between agents
- the notion of *immediacy* or distance between agents, and this affecting their influence upon one another
- an allowance for an agent to be influenced by a *number* of other agents simultaneously

Existing models often incorporate one or more of these concepts, explicitly or otherwise. We will examine these models throughout this chapter. In addition to their use of aspects of Social Impact Theory, we will also assess the suitability of these studies for application to models of *culture*. Cultural models are usually distinguished by agents having a set of discrete traits; 'opinion-based' models tend to only study a singular binary or continuous value within agents [34]. Therefore, we are also looking for models that use **multiple features**, each with **multiple possible discrete trait values**. A summary of these requirements, and whether they are present in existing models, is given at the end of the chapter in Table 2.1.

A general theme in 'social influence' is the tendency for agents to become more alike [34], through the copying and transmission of various attributes. This is usually referred to as *homophily*, and is commonly observed [103]. Many agent-based models tend to be characterised as either opinion-formation models or models of cultural dissemination, but dynamics and characteristics of 'cultural' or 'opinion-based' models will often overlap. Although our model introduced in this thesis (Chapter 3 and extended in chapters 5, 6) is one of cultural spread (i.e., multiple discrete traits), in this section we give a brief overview of some of the more fundamental opinion-based models as some may have dynamics similar to those we will later use. Cultural models are described in their own section, 2.3.

Voter model

The voter model [76] is a particularly simple model of social influence. Each agent is given a single binary opinion, usually assigned randomly at initialisation. At each iteration, a random agent is activated, and one of its neighbours is randomly chosen. The activated agent simply copies the opinion of its neighbour without any other considerations, and the sites converge to a single opinion. While simplistic, the model and its derivatives display behaviours familiar to physicists, such as phase transitions and other dynamics analogous to those of ferromagnetic models [34].

The introduction of 'zealots' to the voter model causes changes to the dynamics. In simple topologies, a single zealot will eventually convert all others to their position [109]. Multiple zealots with conflicting opinions can cause fluctuating steady states [110].

While influential, the voter model lacks some of the key features we need to answer the research questions posed in Chapter 1, and to apply Social Impact Theory to cultural contagion. The voter model features only a single feature of binary value. The notion of *strength* is absent in the base model, influence is only dyadic rather than from a greater *number* of sources, and potential sources of influence are only those most near (see Table 2.1).

Majority Rule

Inspired by public debates, the majority rule model [64] expands on the voter model by allowing multiple sources of influence. The means by which this occurs is on every iteration, a 'discussion group' r is drawn from the entire population. All members within this group will adopt the majority opinion within the group. The size of the group may vary on each iteration, from a given distribution. Where the size of the group is even, deadlock can occur when the number of 'votes' for each opinion is evenly split. Galem frames the binary opinions as 'reform' and 'status-quo'; on a deadlock, the group will choose the status quo opinion. This gives a slight bias to one of the potential opinions. Where group sizes are odd, the majority opinion (whichever it is) will eventually win out and gain consensus in the population. Where group sizes are even, the 'status-quo' opinion will eventually dominate, even if the agents who hold it start as a minority.

The use of a 'discussion group' is analogous to *number* from Social Impact Theory, but the other limitations present in the base voter model still remain if we are to model Social Impact Theory in a cultural model.

A Dynamic Theory of Social Impact

In 1990, Latané worked with Nowak and Szamrej to develop an agent-based model of his social impact theory [122]. As social impact theory is fundamental to this thesis, we look at this model in a little more detail. Its agents each held a single binary opinion (randomly assigned at initialisation), and were placed in cellular grids. Each agent was also given a random value for 'persuasiveness', and a random value for 'supportiveness'; these two variables represent the *strength* of social impact. *Persuasiveness* gives the strength of the agent when influencing those others with an opposing opinion to change it; *supportiveness* gives the strength of the agent when influencing those of the same opinion to resist change. *Immediacy* is based on the Euclidean distance of the influencing agent; applied to the strength using an inverse-square rule. The fact that a target is influenced simultaneously by multiple persuasive and supportive sources gives the *number* element of social impact. This is moderated however, by the 'psychosocial' law of social impact described in Section 2.2.1; these 'diminishing returns' are achieved by by multiplying influence by the square root of the number of influencers.

The authors determine the persuasive (i.e. opinion-changing) impact on an agent as:

$$\hat{l}_p = N_o^{\frac{1}{2}}[\sum(p_i/d_i^2)/N_o]$$

and the supportive (i.e. opinion-affirming) impact on an agent as:

$$\hat{l}_s = N_s^{\frac{1}{2}}[\sum(s_i/d_i^2)/N_s]$$

where N_o and N_s represent the number of agents with opposing and shared views respectively; p_i or s_i the persuasiveness or supportiveness of the source i; d_i the distance between source i and the target.

On each step of the simulation, the rules above are applied to determine whether an agent will change or hold its opinion. The simulation tends to reach a polarised state at equilibrium. One opinion usually forms a clear majority, with clusters of the minority opinion pushed to the margins of the cellular grid.

While the most obvious example of a computational model of Social Impact Theory, this dynamic model of social impact has limitations in regard to potential application to culture models (see Table 2.1). As with voter models, Nowak et al. use a single binary feature, whereas culture models are typically multifaceted. Strength is assigned randomly, rather than taking into account homophily or other potential real-life effects on strength of influence. Influence is from multiple sources, but this includes the entire population; culture models are usually based on more localised influence. However, we adapt the inverse-square method of diminishing influence over distance in our model in Chapter 6.

Continuous opinion models

Probably the most well known opinion model with continuous values is that of Guillaume Deffuant et al. [49]. Here, agents are placed in a graph and may interact with their neighbours. Each agent is given a random starting opinion somewhere between 0 and 1. These boundaries are perhaps the 'extremes' of an opinion, with moderate views lying between. At each iteration, a randomly chosen agent will interact with a randomly chosen neighbour, and their opinions may move closer to each other's.

A threshold to discussion is often set: if the two agents' opinions are too far apart, they will not interact; this barrier to interaction is referred to as *'bounded confidence'*. The amount by which agents move their opinions toward each other is determined by a convergence parameter. There is an element of homophily; agents close enough in opinion becoming more alike, but those too far apart not interacting and perhaps becoming polarised. The bounded confidence threshold determines how many final clusters of opinions will emerge. Where the threshold is large, all agents converge to a single opinion. Where the threshold is smaller, clusters of different opinions may form.

Deffuant et al. have made a number of extensions to the base model, including the introduction of 'extremists' [48]. Here, the uncertainty of agents (i.e. the threshold to interaction) is influenced by others as well as the opinion. Furthermore, an agent's influence on others is increased by low levels of uncertainty - an 'extremist' will be less uncertain in their views. Where there was a high level of uncertainty in the general population, extremists tended to succeed in winning over others to their views. In [2], Amblard and Deffuant examine the effects of network topology (particularly lattices and small-world networks) on this extremism, finding that at low connectivity extremists could spread their opinion more than in high connectivity scenarios. Deffuant also found that convergence around extremist opinions could be made possible by the isolation of moderates [47].

Hegselmann and Krause [74] offer a similar approach, but in contrast to the Deffuant model they incorporate multiple influences. Instead of being influenced by one neighbour at an iteration, an agent adopts the average opinion of *all* its compatible neighbours. The additional cognitive expense of determining this average opinion makes time-to-convergence generally slower than in the Deffuant model.

The use of continuous rather than binary values allows for greater variance in the number and composition of groups formed. However, this single continuous opinion is again quite different to the multiple features of distinct traits used in culture models (see Table 2.1). Some form of homophily is present, but this serves more to block interaction between dissimilar agents rather than as a form of *'strength'* of influence. The Hegselmann and Krause model allows for multiple sources of influence (the *'number'* in Social Impact Theory), but influence over *distance* is absent.

2.2.3 Social contagion

The copying of traits and opinions from one agent to another also bears some similarity to models of *contagion*. Models of contagion are often epidemiological in nature, such as agent-based SIR models. In their simplest form, an SIR model [81] simulates the spread of a trait (for example, a disease one either has or does not) throughout a population. Each member of the population has a status from **S**usceptible (does not have the disease or trait but is able to receive it), **I**nfected (has trait), or **R**ecovered (no longer has trait, and may or may not be immune from receiving it again). These models and their derivatives have been re-purposed for scenarios beyond those of infectious diseases, such as the spread of news [79], rumours [43], financial crises [51], corruption [118], memes [157], scientific ideas [66], and more besides.

Studies applying network analysis to contagion are myriad [128]. In the context of epidemiology, applying network science principles gained momentum with the study of the spread of sexually transmitted diseases in the 1980s [95]. Even where contagion and graph theory are not explicitly mentioned, earlier work hinted at the importance of short paths of transmission; perhaps most famously in Travers and Milgram's 1969 'small-world' experiment that popularised the notion of 'six degrees of separation' [152]. Granovetter further advanced the importance of global path-length and how it is typically the 'weaker' ties of social acquaintances that contribute to the small-world phenomena [69]. Strong ties in a network tend to be clustered, and provide a degree of local redundancy; where information spreads through a network, a node's closest contacts will tend to also be close to each other, and receive the same information almost simultaneously. Weak ties on the other hand are often *long*, joining otherwise disparate communities. Watts and Strogatz found that it only takes a relatively small number of these ‘long’ ties acting as shortcuts to give a clustered network a low average path-length, similar to that of a random graph [160]. However, unlike a random graph, a Watts-Strogatz graph can retain a high degree of clustering while the graph diameter becomes short: a key characteristic of a ‘small-world’ network. These types of network properties could have significance for the behaviour of our proposed model, and will be key to answering research question RQ4 (see Chapter 4).

Although a highly contagious disease or rumour may jump across such long weak links easily, the spread of behaviours as a ‘viral’ sensation is disputed. While a virus may spread alarmingly easily, measures to prevent it may meet resistance [127]. In some cases, non-adherence to measures was caused or exacerbated by peer pressure and rumours of side effects [155]. There are instances where doctors are reluctant to adopt a new drug unless their colleagues are using it [41], or activists are more likely to join a social movement where it aligns with their other social ties [102]. Damon Centola argues that where social behaviours are concerned, the threshold to adoption is often higher, and therefore diffusion requires multiple sources of influence - a ‘*complex contagion*’ [36]. As our model intends to incorporate *number* from Social Impact Theory as simultaneous multiple influences, there could be links between the theory of complex contagions and our work. In particular, simple and complex contagions are key to answering research question RQ3, so we describe this theory in more detail below.

Simple and Complex Contagions

Damon Centola makes a distinction between *simple* and *complex* contagions. The definitive difference is the *minimum* number of contacts needed for a trait to spread. In simple contagions a single interaction with a single contact may be sufficient for transmission. This is the case for many infectious diseases; while repeated interactions may increase the chances of contagion, the minimum possible contact required for transmission is still a single interaction. Travellers forced into close proximity on a metro train, who may never have met and will likely never see each other again, may nevertheless catch a common cold from one another. As we are unfortunately all too aware, even pressing the same elevator button as an infected person, then touching one's face, may be sufficient contact for the transmission of a potentially deadly virus.

Complex contagions however, require multiple sources for a trait to spread. It is argued that the spread of human behaviours often require social reinforcement from multiple sources. This is often because the adoption of a new trait or behaviour incurs some form of cost, be it financial, social, or otherwise. In communications technologies, adoption will likely depend on whether the potential user perceives the benefit of the technology to exceed its cost; yet the value derived from the service is directly dependent on the number of its users (a 'network effect') [94]. One is unlikely to purchase a fax machine or mobile phone without other compatible devices to communicate with. Multiple trusted sources adopting a behaviour increases its credibility, such as cases where doctors are more likely to trust the use of an innovation once their colleagues are seen to have adopted it. The risk of damaging one's reputation by adopting a new behaviour is reduced when others make the same decision [36]. The idea of a 'critical mass' is common across varied social contexts, from seminar attendance to collective action [100, 138]. Where contagions are complex this can impact widely held assumptions regarding how behaviours spread 'virally' or across 'weak links' within small-world networks.

A comparison of these contagion dynamics with those of our model is the focus of research question RQ3. Consequently specific mechanisms and local dynamics are described in more detail in Section 4.1.1.

2.3 Culture

The word 'culture' has many different definitions and comes with various connotations. Kroeber and Kluckhohn listed over 150 definitions in 1952 [87]; in 2006 Baldwin et al.

added 300 more [13]. The majority at least frame culture as something we learn from each other. For the purposes of this thesis we shall use the same very broad definition that Robert Axelrod uses in his 1997 work *The dissemination of culture*: the *'set of individual attributes that are subject to social influence'* [6].

Axelrod goes on to state that cultural diffusionists tend to treat culture as a set of distinct traits, each of which can be passed on. Many computational models work to the same principle; the general perspective is that this set of categorical and non-ordinal traits distinguishes culture models from opinion formation models [34]. These traits are usually abstract; they are not attached to any real-world meaning or significance. What is important is that they may spread, and global states of distinct cultures of different numbers and sizes emerge. A large number of subsequent cultural models are based on Axelrod's original, including the model we develop over the course of this thesis. Accordingly, we describe his model in more depth below.

2.3.1 Axelrod's model of Cultural Dissemination

As we have mentioned above, Robert Axelrod's influential paper *The dissemination of culture* [6] modelled culture as a set of discrete traits; a representation that has been adopted by many subsequent works (see Section 2.3.2). Where previous models of social influence (such as the voter model and its derivatives), tended to treat features individually, Axelrod took into account the interaction between features. This is key to the homophilic copying within the model; the more existing features agents have in common, the more likely they are to copy additional traits from each other. Axelrod set out to investigate the apparent paradox of cultural convergence and diversity. If people tend to associate with those more similar and adopt their behaviours and norms, why then do cultural differences continue to exist? In a similar result to Schelling's earlier work on residential segregation [137], it was found that even a small preference for interacting with those more similar at a local level can create a global state of polarisation. Different cultures continue to exist, but they are disparate and essentially disconnected from each other.

In Axelrod's model, agents are each assigned a list of cultural features (or 'dimensions of culture'). Each feature takes a value from a set of discrete cultural traits. For example, one feature may represent a piece of clothing, and each possible trait value may represent a style for that item. Agents in this model do not necessarily represent individuals, but could represent a 'cultural site'; this could be an individual, a village, a nation, or any entity with a cultural identity. Expressed more formally, each agent j holds a feature vector $(\sigma_1^j, \sigma_2^j, \ldots, \sigma_F^j)$ where there are F features[1]. Each feature is populated by one value from a set of q traits. Let T denote the set of traits. In Axelrod's work and subsequent extensions, these traits are typically denoted by numerals; however it is important to remember that they are discrete, and the value 1 is no more similar to 2 as it is to 9. Any mutually exclusive symbols could be used to denote traits, but alphanumeric characters are usually most convenient; thus a set T of $q = 10$ traits will almost always be represented by the integers 0 to 9. In this way an agent's culture is defined by its features; an example where $F = 5, q = 10$ might be $(0, 7, 4, 8, 1)$.

In the original work agents are assigned random trait values and placed in a simple square lattice, often a 10x10 grid of 100 agents, reminiscent of cellular automata. An agent may interact with those to its north, east, south or west — i.e. a 'Manhattan Distance' of 1. This grid is bounded and non-toroidal; the agent at the extreme north-west square of the grid only has neighbours to its south and east.

Trait-spreading interactions in Axelrod's model are based on *homophily*, the widely observed [103] tendency for individuals to associate with those most similar to themselves. On each iteration a random agent i is activated. It chooses a random neighbour, j, to interact with. On interaction, the existing similarity between i and j is calculated:

[1] Typically a set is denoted by capitalisation, and its size by $|F|$, but in the interests of consistency with other extensions of Axelrod's work, we will always use F to refer to the *size* of the feature vector rather than the vector itself.

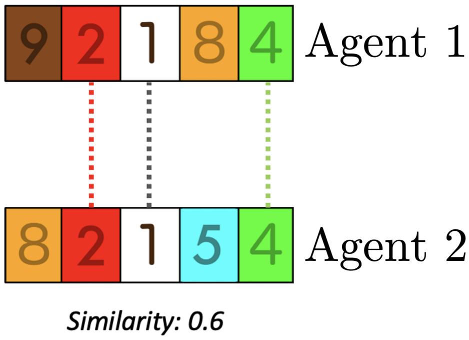

Figure 2.2: Trait copying in Axelrod's model: These two agents share 3/5 features in common (features 2, 3 and 5), giving a similarity of 0.6. Should Agent 1 be activated, they therefore have a 60% chance of copying from Agent 2. If this occurs, a random dissimilar feature will be chosen (i.e. feature 1 or 4 in this case), and Agent 1 will adopt the trait agent 2 holds in that feature. For example, Agent 1 may change its 4th feature to be trait 5, to match Agent 2.

$$sim_{i,j} = \frac{1}{F}\sum_{k=1}^{F}\delta_{\sigma_k^i,\sigma_k^j}$$

where $\delta_{\sigma_k^i,\sigma_k^j}$ is Kronecker's delta. Therefore $\delta_{\sigma_k^i,\sigma_k^j} = 1$ only when the k^{th} feature of i and j are the same. Pseudo-code for calculating this existing similarity is given in Algorithm 2.1. This calculation is based on the number of features the agents have in common, and this similarity dictates the likelihood that i will copy a trait from j. Where i and j have two out of five features in common, the probability of copying is $\frac{2}{5}$, or 0.4. Should this occur, i will adopt a trait from j for one of the remaining dissimilar features. An example of this copying process is given in Figure 2.2, pseudo-code in Algorithm 2.2.

In this way, similar agents become more similar, driving convergence. There may be instances where two neighbours have traits in common, but lose this commonality as one is influenced in a different direction by their other neighbours. However, in general, a small amount of similarity often leads to complete similarity. Even if two adjacent agents have low similarity and probability of copying, given enough time they will

Algorithm 2.1 calculateSimilarity(i,j), based on the number of features the two compared agents have in common.

Require: two agents, i and j
Require: the number of features, F

```
similarity ← 0
for k ← 1 to F do
    if σ_k(i) = σ_k(j) then
        similarity ← similarity + 1
    end if
end for
return similarity/F
```

copy from either each other or their other neighbours. The simulation can only reach a state of stability where the possibility of copying traits is absent from all relationships between agents in the population. There are two circumstances where an agent has zero chance of copying from a neighbour: it has exactly the same culture, or it is completely dissimilar. Convergence and polarisation are inevitable and inescapable. An algorithm for determining a stop condition is given in Algorithm 2.3; note that this occurs when all agents are adjacent to either completely similar or completely dissimilar neighbours.

Axelrod's principal measure for this convergence was the number of *cultural regions*, each being a set of contiguous sites with identical culture. His work also defines *cultural zones*, areas containing agents with 'compatible' culture; i.e. at least one feature in common. Cultural zones eventually become cultural regions of identical agents. The number of these cultural regions that the stabilised simulation can support is largely affected by the number of features F and traits q. An increase in the size of the feature vector F gives more chance of agents having at least one trait in common, driving the copying behaviour that eventually leads to homogeneous cultures. An increase in the number of possible trait values q for each feature increases the chances that agents will have nothing in common, allowing different cultures to persist. When trait values are

Algorithm 2.2 Interactions in Axelrod's Model

Require: the number of features, F

Require: a global set of agents A, each containing a set of F features

$stabilised \leftarrow$ **false**
while not $stabilised$ **do**
    Select random agent i from A where $|N(i)| > 0$
    Select random agent j where $j \in N(i)$
    $similarity \leftarrow calculateSimilarity(i, j)$ // see Algorithm 2.1 for $calculateSimilarity(i, j)$
    $randomInt \leftarrow$ a random integer between 1 and F inclusive
    if $F > (similarity \times F) \geq randomInt$ **then** // '$F >$' ensures agents not already identical
        Select random feature k where $0 < k \leq F$ and $\sigma_k(i) \neq \sigma_k(j)$
        $\sigma_k(i) \leftarrow \sigma_k(j)$ // i copies a trait from j
    end if
    $stabilised \leftarrow isStabilised()$ // see Algorithm 2.3
end while

Algorithm 2.3 Detecting stabilisation in Axelrod's model. Stabilisation can only occur where all agents have either complete similarity or complete dissimilarity with those they may interact with.

Require: a global set of agents A, each containing a set of F features

for all agent $i \in A$ **do**
    for all neighbour j of agent i **do**
        $similarity \leftarrow calculateSimilarity(i, j)$
        if $0 < similarity < 1$ **then** // agents are neither identical nor completely dissimilar
            return false
        end if
    end for
end for
return true

```
92491 92491 92491 14805 14805 14805 92491 92491 92491 92491
92491 92491 92491 14805 14805 14805 92491 92491 92491 92491
92491 92491 92491 14805 14805 92491 92491 92491 92491 92491
92491 92491 92491 14805 92491 92491 92491 92491 92491 92491
92491 92491 92491 14805 92491 92491 92491 92491 92491 92491
92491 92491 92491 92491 92491 92491 92491 92491 92491 92491
92491 92491 92491 92491 44269 92491 92491 92491 92491 92491
92491 92491 92491 92491 92491 92491 92491 92491 92491 92491
92491 92491 92491 92491 92491 92491 92491 92491 92491 92491
92491 92491 92491 92491 92491 92491 92491 92491 92491 92491

---------------------------------------------------------------------95700
 Num different Cultures: 3
 Stabilised near iteration: 95700
```

Figure 2.3: A typical end state of Axelrod's model, where $F = 5, q = 10$. Here we have recreated Axelrod's model, and it has stabilised with 3 cultural regions. Note that each agent is either identical or completely dissimilar to its neighbours; the cultures are completely polarised. Interestingly, the cultures marked in blue and red actually do have one compatible feature (feature 2), but they are insulated from each other by the dominant '92491' culture. Had the blue and red cultures been geographically adjacent, one would have assimilated the other.

assigned randomly, the probability that two neighbours will start with all features in common is $(\frac{1}{q})^F$, that they will have no features in common $(1 - \frac{1}{q})^F$, and that they will have at least one feature in common $1 - (1 - \frac{1}{q})^F$ (Figure 2.4). These starting probabilities give an indication of how likely it is that multiple cultures will exist when the stabilisation is run. However, the final number of cultural regions is not a direct translation from these starting probabilities. When the simulation is complete, at low values of q multiple cultures cannot be sustained at all; at high values of q the population will consist of many very small regions. Physicists have characterised this as a phase transition [35], in a similar fashion to the abrupt transitions between basic states of matter under temperature changes.

Axelrod's model continues to be a popular basis for simulations of social and cultural influence. As mentioned above, and in [34], subsequent computational models tend to follow Axelrod in treating culture as multi-faceted.

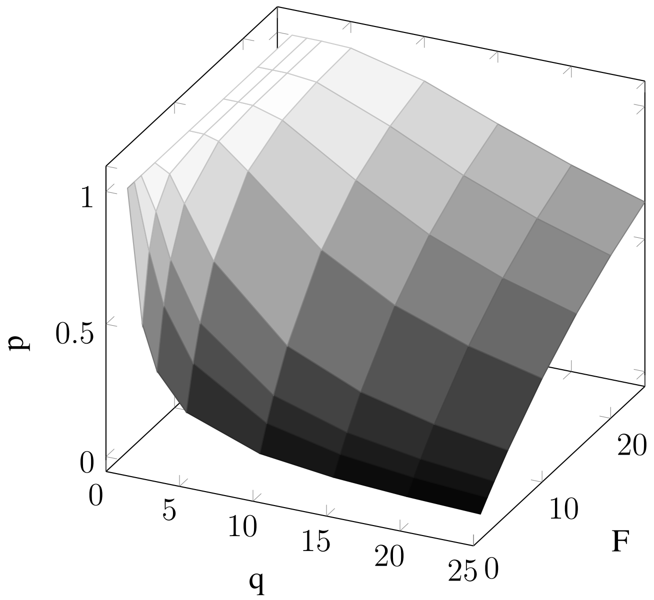

Figure 2.4: Starting probabilities (p) of at least one feature in common for different values of q and F.

2.3.2 Extensions to Axelrod's model

There are numerous extensions and comments on Axelrod's model, from drawing analogies with similar phenomena in physics [35], to applications in optimisation problems [60].

Many papers examine the effects of different network topologies. In [85], the model's effects in small-world and scale-free networks were examined, while [143] utilised Erdős-Rényi random graphs. Battiston et al. [16] take a different approach involving multiplex networks where an agent may have different contacts for each feature. For example, one may be able to influence another individual on sport but not politics.

Mobility of agents has also been considered. Gracia-Lázaro et al. [68] do away with the usual network structures and instead place the agents in a world more akin to that of Schelling's residential segregation model [137], where agents may move to empty spaces in an Axelrod-Schelling hybrid. In [37], Centola et al. allowed a form of movement within a network: agents dynamically 'rewiring' their edges in a network which co-evolves alongside the cultures. In this model, where an agent has zero similarity with one of its neighbours, it may abandon contact with it and instead create a link to another randomly chosen agent. At certain values of q (number of traits), the network will disconnect into disparate sub-graphs, each of a single culture.

Similar to introducing ‘zealots’ or ‘extremists’ to opinion formation models [48, 109, 110], Singh et al. [143] add ‘committed agents’ to the Axelrod model, finding that beyond a critical number of committed agents convergence time grows logarithmically. Both Shibanai et al. [140] and González-Avella et al. [67] have added the effect of mass media to the model, by modelling a mass media message as a ‘global other’ with which agents may interact. In [67], on each iteration there is a probability B that an agent may be influenced by the ‘mass media’ rather than one of its local neighbours. Counter-intuitively, high values of B can induce multiple cultures rather than convergence to a monoculture, due to increased polarisation with those agents who do not share traits with the media message.

The notion of ‘cultural drift’ is borrowed from evolutionary biology’s ‘genetic drift’; minor errors or fluctuations that may lead to broader change. Klemm et al. [84] introduce cultural drift to Axelrod’s model as *noise*: there is a small probability that an agent will change a trait randomly and not through social influence. These small perturbations can ‘shake’ Axelrod’s simulation out of a state of multiple polarised cultures and into a global monoculture. Flache and Macy [59] demonstrate that the fragility of cultural diversity in the face of these random perturbations can be made more robust by influence being from multiple sources rather than one, a method thematically close to Social Impact Theory’s *number*. Consequently we examine this model in further detail below.

Flache and Macy’s model of Social Influence

In [59], Flache and Macy expand Axelrod’s model to include *noise*, *selection error*, and what they refer to as *social* rather than *interpersonal* influence. The authors characterise the dyadic interactions of Axelrod's model as 'interpersonal', and simultaneous multiple influences as 'social'.

At each activation, the node in question i will add each of its neighbours j to an influential subset S with probability p_{ij}. This is somewhat similar to the 'Majority Rule' voter model of social influence (see Section 2.2.2), however the group is drawn only from i's neighbours, and only i will be updated rather than the whole group. This subset S will contain only the nodes who may have influence over i during this iteration. The probability of each neighbour j of i being added to this subset is denoted as p_{ij} by Flache and Macy; this measure is either always 1 if homophily is disabled, otherwise it is identical to the similarity measure used in Axelrod's original model. However, Flache and Macy also use selection error; with probability r', the decision on whether to include or exclude a neighbour j from group S will be *reversed*. In addition to selection error, noise is also added; at intervals in the simulation an agent may change one of its features to a randomly chosen trait.

When S has been formed, and assuming it is not empty, a feature k is randomly chosen for updating[2]. For each possible trait α of k the number of 'votes' α receives is counted; that is, each member of S that holds that trait on that feature contributes a vote toward it being adopted. The trait with the maximum votes, α', is adopted for that feature k by agent i.

Despite the disruptive dynamics of selection error and noise, the multiple social influences allowed in the model offer greater robustness against random perturbations pushing the simulation into a single monoculture.

[2] In their paper, Flache and Macy use the symbols F, Q, f and q to refer respectively to the number of features, number of traits, feature in question, and trait in question. For consistency with Axelrod's model we use F and q throughout this thesis to denote the number of features and number of traits. Specific features and traits are indicated by k and α respectively, as used in our trait-scoring formula in Chapter 3.

2.4 Summary

Numerous approaches and models have been developed to offer potential explanations for aspects of how human beings influence each other, and spread behaviours, values, opinions, and norms. Many models seek inspiration from physics for their mechanisms, and are often analysed in comparison to physical phenomena. Few use social impact theory deliberately or explicitly, although some have processes which could be thought of as analogous to some aspects of social impact theory; these are compared in Table 2.1. Nowak et al. [122] have contributed a well known model of social impact theory characterised as a physical force, but this is used as an opinion formation model with only a single binary opinion being influenced rather than the multiple discrete traits typical of a culture model. Although Flache and Macy's culture model [59] can be used as one of multiple influences, social impact is not the focus of the model. Consequently there remains a gap in the literature for a systematic understanding of cultural contagion using the mechanisms on which social impact theory is based, namely the *strength, number,* and *distance* of influences. This is the focus of our thesis, which begins by introducing a new modelling approach for this purpose in Chapter 3.

Model	Features F	Traits q	Potential for utilising social impact		
			Strength	**Number**	**Distance**
Voter [76]	1	Binary	Not present, agents always copy from neighbours	Single neighbour	Local neighbours only
Majority [64]	1	Binary	Agents have equal strength of influence	Group: randomly created group drawn from whole population	Fixed: whole population
Dynamic social impact [122]	1	Binary	Random distribution of strengths among agents	Whole population (cellular automaton)	Fixed: Whole population, but with inverse-square law of diminishing influence over distance
Deffuant [49]	1	Continuous	Agents that are too dissimilar are blocked from interacting, otherwise agents have equal strength of influence	Single neighbour	Local neighbours only
Hegselmann & Krause [74]	1	Continuous	Agents that are too dissimilar are blocked from interacting, otherwise agents have equal strength of influence	Group: entire local neighbourhood	Local neighbours only
Axelrod [6]	Multiple	Discrete, multiple	Homophilic copying	Single neighbour	Local neighbours only*
Flache & Macy [59]	Multiple	Discrete, multiple	Not explicit; homophily influences *number*	Group: stochastically created group drawn from local neighbourhood, probability of membership based on homophily	Fixed: close neighbours only*
Our proposed model	*Multiple*	*Discrete, multiple*	*Always considers homophily, and optionally status*	*Group: all nearby agents within distance d*	*Variable distance d. Multiple possible methods of diminishing influence over distance*

Table 2.1: Comparison of models of social influence. We have identified aspects of these models which could be seen as analogous to the social forces within Social Impact Theory, even if the model was not designed as such. *Both Axelrod and Flache and Macy describe distance as a simple neighbourhood extending out from the agent. In a network graph this is more an increase in connectivity than distance. This is discussed in further detail in Section 6.3.

Chapter 3

An agent-based model of social impact based cultural dissemination

In this chapter, we introduce our model which will act as the basis of experiments throughout this thesis. Extending Axelrod's *Dissemination of Culture*, we use Latané's *Social Impact Theory* as the basis of social influence within a culture model. In particular, this chapter focuses on expanding the influence upon an individual to include multiple sources of local influence; a mechanism that can give markedly different results when compared to dyadic copying. We also outline alternative metrics for examining the state of cultural models on network graphs.

The work outlined in this chapter has previously been published in part in [112]. The measures of polarised, homogenised, and mixed links outlined in Section 3.3.2 have been published in [113]. This chapter supports research questions RQ1 and RQ2.

3.1 Introduction

In his work on cultural dissemination, Axelrod focuses on cultural 'sites', which may or may not necessarily represent individuals. Indeed, Axelrod writes of implications for state formation and transnational integration as well as social influence between individuals [6]. We also aim to keep our model abstract and generalised, such that the agents may potentially represent any entity (individual or group) with mutable attributes, and the meaning of those attributes is not defined. What is more important is that, as in Axelrod's work, the attributes are multiplex, mutable, and transmittable. However, when it comes to the mechanisms of this transmission, our work extends Axelrod's by embedding Social Impact Theory [88] as a process of social influence. While this theory has often been used as a qualitative description of how an individual may be influenced by others, we attempt to use its mechanisms as alternatives to the dyadic copying of traits within the Axelrod model.

3.1.1 Basis for our model

The extended mechanisms of social influence within our model are drawn from Latané's social impact theory [88], which has previously been adopted to emulate larger scale effects of a range of different social and psychological processes [25, 78, 119]. Described in more detail in Section 2.2.1, the theory characterises influence based on three variables: *strength* (denoted S), *number* (denoted N), and *immediacy* (denoted I). 'Strength' reflects the intensity of a source upon the target being influenced, determined by factors such as status, resources, in-group membership, power. 'Number' is simply the number of influences upon the target. 'Immediacy' represents the proximity in temporal, physical or social space and absence of barriers or filters. Influence in Latané's initial work is given as a multiplication of these factors, $Impact = f(SIN)$.

In basing the probability of copying on an agent's existing similarity with its influencer, Axelrod's model has in effect already included an element of *strength*. The propensity of individuals to behave more favourably toward those of greater similarity is a long-standing one [90, 103]. We retain this element of homophily in our model as an initial basis for the strength of influence. However we allow for the possibility of later extending strength to incorporate other aspects, such as status, which we include in Chapter 5.

A target being influenced by several sources of influence is notably absent in Axelrod's model. We seek to add this *number* aspect, moving away from dyadic interactions to ones where an agent is influenced by multiple sources simultaneously. An individual's closest ties acting upon them simultaneously can in a sense be thought of as 'peer pressure' [112]. This addition is the main focus of this chapter, and also allows our model to be used as a tool for examining the effects of dyadic versus multiple influence in other networks and scenarios (chapters 4, 5, 7).

Latané's theory also contains a 'psychosocial law', whereby each additional source of influence contributes a little less (see Section 2.2.1). In our model influence will be divided by the number of neighbours, giving an effect similar to that outlined in the original theory [88] and the model of Dynamic Social Impact Theory [122].

The third social force of Latané's social impact is that of 'Immediacy' or *distance*. We defer the consideration of this factor until later in this thesis (Chapter 6); instead we first seek to isolate and examine the effects of *number* of influences.

3.1.2 Other models of multiple influence

In Chapter 2, we described previous models that have some element of multiple influence, although few explicitly set out to model Social Impact Theory. The 'Majority Rule' [64] (Section 2.2.2) extension of the voter model allows potential influence from up to the whole population, but these influences are chosen stochastically rather than having a basis in the agent's local network. Hegselmann and Krause [74] (Section 2.2.2) narrow their multiple influences to a local neighbourhood, but theirs is a model of continuous values. Both of the models above are usually classified as opinion formation models, with a single feature rather than the vector of features with discrete values we associate with culture models.

The model of Nowak et al. [122] (Section 2.2.2) *is* explicitly one of Social Impact Theory, however it is also framed as an opinion formation model with a single binary feature. Also, it is based on cellular automata with agents influenced by the entire population (albeit with a diminishing effect over Euclidean distance); we are looking to create a model of cultural diffusion on network graphs to allow variations in network structure beyond simple lattices.

Flache and Macy's model of social influence [59] also contains a mechanism of multiple influences, but their model focuses on noise and selection error, which are beyond the scope of our current work. However, as it is a cultural model of multiple features, we will use an adaptation of it (with homophily always present but zero noise and zero selection error) as a benchmark for testing our model.

3.2 The Model

Let A denote the set of agents. Rather than restrict attention to a cellular grid, agents form the vertices in an arbitrary undirected network graph, and are connected to their immediate neighbourhood. Thus where a Graph G would normally be defined as a pair of sets, $G = (V, E)$, here the set of vertices is A. We consider one graph at a time, and we denote the neighbourhood of a node i as $N(i)$ (i.e., a set of vertices adjacent to i). The equivalent network graph to the grids used in Axelrod's paper would be a basic grid lattice, but any connected network can be used. Initially, edges are undirected and unweighted. Neighbours may interact with each other in either direction with equal weight. Similarities could have been used as edge weights, but we reserve edge weights for other usages in later chapters. Also, the similarities between agents will change and need to be recalculated frequently.

Each agent i holds a feature vector $(\sigma_1^i, \sigma_2^i, \ldots, \sigma_F^i)$ where there are F features[1] and each feature is populated by one of q traits. Let T denote the set of traits. As with

[1]As previously mentioned when describing Axelrod's model in Section 2.3.1, typically a set is denoted by capitalisation, and its size by $|F|$, but in the interests of consistency with other extensions of Axelrod's work, we will always use F to refer to the *size* of the feature vector rather than the vector itself.

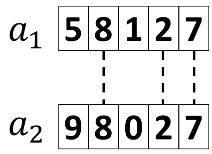

Figure 3.1: An example of two agents (a_1 and a_2) with their associated feature vectors. Assuming $F = 5$ and $q = 10$, using traits denoted 0-9, these agents have three common features (positions 2, 4 and 5). Their similarity $sim_{i,j} = 0.6$. Figure taken from [112], ©2019 IEEE.

Axelrod's work, it is important to remember that although often depicted by numerals, traits are discrete, non-ordinal and mutually exclusive. The value 1 is no more similar to 2 as it is to 9.

As in Axelrod's model, the similarity, sim_{ij} between neighbours i and j is calculated:

$$sim_{i,j} = \frac{1}{F} \sum_{k=1}^{F} \delta_{\sigma_k^i, \sigma_k^j}$$

where $\delta_{\sigma_k^i, \sigma_k^j}$ is Kronecker's delta. This similarity will form the basis of the notion of *strength* in our model; the more similar the influencing node, the stronger its influence. An example is given in Figure 3.1, pseudo-code in 2.1.

Similarly to Axelrod's model, at each iteration a random agent is activated. This retains a stochastic element in the model, as the mechanism for copying is otherwise deterministic. The random order in which agents are activated can produce different local results for simulations even where starting conditions are identical.

On activation, an agent is influenced by all of its neighbours simultaneously. This composite influence from the immediate neighbourhood constitutes the *number* aspect of our model, as influenced by the number of sources of influence in Latané's social impact theory. For each feature k (where $k = 1$ to F) of the activated agent i, each possible trait $\alpha \in T$ is scored. This trait-score is calculated by examining the similarity of i to every neighbour which holds the trait α for feature k. These similarities are summed, and the result is the score for that particular trait. This process is repeated for every possible trait of feature k; the trait $\alpha' \in T$ which scores highest is selected as the new value of k for the agent i. Where the current trait value of feature k is the highest or joint-highest scoring trait, then this value is retained. Should there be multiple traits holding the same highest score, and none of these traits are the current value of k, then one of these highest-scoring traits is adopted at random.

More formally, the trait scoring formula is:

$$ts_{\alpha,k,i} = \sum_{j \in N(i)} sim_{ij} \delta_{\sigma^j_k,\alpha}$$

where $ts_{\alpha,k,i}$ denotes the trait-score of trait α for feature k of agent i; $N(i)$ is the set of all neighbours of node i; sim_{ij} is the similarity between i and j (see above). $\delta_{\sigma^j_k,\alpha}$ is Kronecker's delta, equal to 1 only when the k^{th} value of j is identical to the trait α being scored. Otherwise, the trait scores zero for j and the process moves on to the next neighbour of i. When all traits have been scored, $\alpha' \in T$ is determined such that $ts_{\alpha',k,i} \geq ts_{\alpha,k,i}$, $\forall \alpha \in T$. Then σ^i_k is set to α. An example of this process is given in Figure 3.2. Pseudo-code is outlined in Algorithm 3.1.

3.2.1 Stabilisation

For comparison with other models, we are largely interested in the state of these simulations at their stable end point. Were a simulation to be halted early, it may present a state not indicative of the cultural states the model dynamics ultimately lead to.

An algorithm for detecting stabilisation in the Axelrod model is given in Algorithm 2.3 (Section 2.3.1). However, for our non-dyadic model we cannot be sure that the polarisation intrinsic to Axelrod's model is as inevitable in ours. The activation process of our model remains random, agents may be called in any order. However, the copying mechanism upon activation is deterministic. For a given time-step, an agent will always

Algorithm 3.1 Interactions in Our Model. Note that this is much simplified, we have excluded code that records data for the purposes of statistics, and trivial calculations that may be language specific (e.g. obtaining the traits with maximum value from the $traitScores$ array). In reality we also configure the check for stabilisation to be at greater intervals rather than on every iteration, for reasons of optimisation.

Require: the number of features, F
Require: the number of traits, q
Require: a global set of agents A, each containing a vector of F features
for all $a \in A$ **do**
  $stable_a \leftarrow false$ // set all agents to stable = false
end for
$stabilised \leftarrow$ **false**
while not $stabilised$ **do**
  $iterationCount \leftarrow iterationCount + 1$
  Select random agent i from A
  $stable_i \leftarrow true$ // Used to determine stabilisation
  for $k = 1$ to F **do** // for each feature
    $traitScores \leftarrow (ts_0, ts_1, ..., ts_{q-1})$
    for $t = 0$ to $q - 1$ **do** // for each trait
      for all $j \in N(i)$ **do** // for each neighbour j
        if $\sigma_k(j) = t$ **then** // agent j has trait in question t for feature k
          $similarity \leftarrow calculateSimilarity(i, j)$ // see Algorithm 2.1
          $traitScores_t \leftarrow traitScores_t + similarity$
        end if
      end for
    end for
    $topTraits \leftarrow \{x : ts_x = max(traitScores)\}$ // all traits with highest score
    if $\sigma_k(i) \notin topTraits$ **then** // if current trait is not in topTraits, it has been outscored
      $adoptTrait(traitScores, i, k)$ // see Algorithm 3.2
    end if
  end for
  $stabilised \leftarrow isStabilised()$ // see Algorithm 3.3
end while

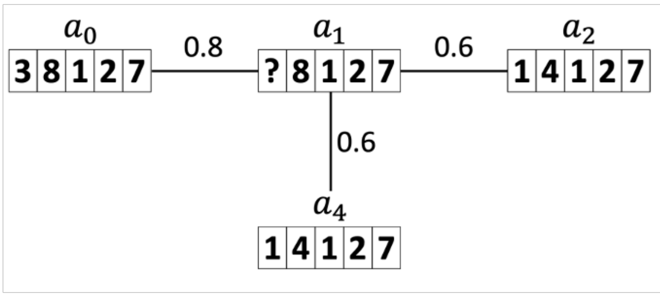

Figure 3.2: Multiple-influence copying based on similarity. Here, where $F = 5, q = 10$, a_1 considers which trait to adopt for feature σ_1. Each possible trait for σ_1 is scored. As traits 0, 2 and 4-9 are not held in feature σ_1 by any agents in the neighbourhood those traits each score zero. The score for trait 1 is calculated as $0.6 + 0.6 = 1.2$; the sum of the similarities to a_1 of the agents which hold this trait (a_2, a_4). Trait 3, held only by a_0, is scored 0.8. Thus despite a_0's greater sole influence over a_1, a trait held by a_2 and a_4 is able to outscore it. Agent a_1 will adopt trait 1 as its first feature.

respond the same way to local conditions. This being the case, we can determine that should every agent in the simulation be given an opportunity to change trait, and none having done so, then no more chances can or will take place (see Algorithm 3.3). The agents have made their choice, and the only circumstances in which they will deviate from that choice is if there is change amongst their local neighbours. Since the whole population has settled and is not changing, the simulation is stable.

3.3 Simulations in square lattices

Initially, we place the agents in a 10×10 regular lattice. Although unlikely to be realistic from the point of view of social networks, this structure allows for direct comparison with Axelrod's model of cultural polarisation [6] and subsequent extensions [35, 68]. As in Axelrod's original study [6], agents are created with randomly assigned starting traits. On each run of the simulation, we are interested in its state at stabilisa-

Algorithm 3.2 adoptTrait() function in our model. Called when one or more traits outscore the current trait for a feature. Where the top score is held by multiple traits, a trait to adopt is chosen at random from those highest scoring traits. When a trait is successfully adopted, all agents are set to potentially unstable (see Algorithm 3.3).

Require: a global set of agents A

Require: a list of highest scoring traits, $traitScores$

Require: an agent i

Require: a feature number k, such that $\sigma_k(i)$ represents the value of the k^{th} feature of i.

if $|topTraits| > 1$ **then** // if multiple traits with highest score
    $adoptedTrait \leftarrow random(topTraits)$
else
    $adoptedTrait \leftarrow topTraits_1$
end if
$\sigma_k(i) \leftarrow adoptedTrait$ // adopt new trait
for all $a \in A$ **do**
    $stable_a \leftarrow false$ // every time a trait is updated, reset every agent to unstable
end for

tion, when no other changes can take place. Descriptions of how this is determined can be found in Section 3.2.1. We examine these end-states over several runs and take the mean values of metrics at the simulation end-point.

3.3.1 Cultures, regions and zones

In Axelrod's initial study [6] and subsequent related works [35, 37], typical measures are the number of cultural regions and zones, and the size of the largest culture. A cultural region is defined as a "set of contiguous sites with an identical culture"; a cultural zone as a "set of contiguous sites, each of which has a neighbour with a 'compatible' culture" [6], i.e. neighbours with at least one feature in common.

Algorithm 3.3 isStabilised() function in our model. We can be certain our model has stabilised if since the last change of trait, every agent has had a chance to copy a trait, and none have since done so. On activation, an agent i is set to stable (see Algorithm 3.1, line 6). If it changes trait, all agents including itself are set to potentially unstable (Algorithm 3.1, line 19). We reach global stability where all agents can be said to be stable since the last trait change.

```
Require: a global set of agents A
for all a ∈ A do
    if stable_a = false then
        return false // agent a has not been activated since last change
    end if
end for
return true // only reachable if all agents have been activated since last successful change
```

In Axelrod's model, the development of cultural zones during a run of the simulation gives an early indication as to how many stable cultural regions will result; at stabilisation the number of zones and regions is the same. This is not usually the case in our model. Polarisation is a condition of stabilisation in Axelrod's model; it can only be said to have stabilised when all agents are either completely similar or completely dissimilar to their immediate neighbours. However, in our model where agents are influenced simultaneously by multiple sources, a stable state can be reached where agents retain some, but not all, of their neighbours' traits. Mixing of adjacent cultural sites can occur without complete convergence or polarisation. This key difference is illustrated in Figure 3.3.

As the number of features F rises in Axelrod's model, the number of stable cultures is reduced. With an increase in the size of the features vector, there is an increased chance of two agents finding at least one feature in common, driving convergence and resulting in a smaller number of cultures - often just one monoculture. Conversely, increasing the number of possible values per feature - the cultural traits, q, reduces the chances of

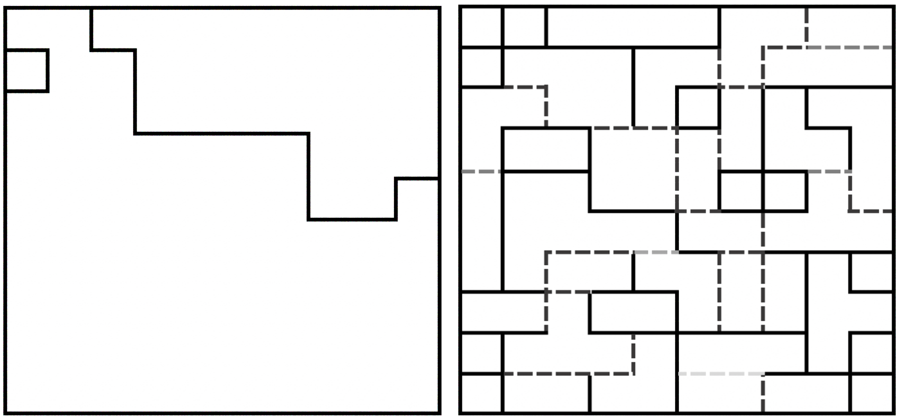

(a) Axelrod's model (b) Social Impact inspired model

Figure 3.3: A unit-square representation of agents positioned on a 10×10 regular lattice. This representation is similar to the 'Map of cultural similarities' used in Axelrod's paper [6]. Solid lines indicate boundaries between culturally dissimilar agents; lighter lines indicate a greater level of similarity between neighbours. The absence of any dividing line indicates complete cultural similarity. Both (3.3a) and (3.3b) were run with the same starting cultures and random seeds, with 100 agents, 5 features and 10 traits. Axelrod's model (3.3a) stabilises at 3 regions of culturally identical agents, yet each region being completely polarised from the next. However, when agents take into account multiple influences simultaneously, stability can occur with mixed cultural zones still existing (3.3b). Figures taken from [112], ©2019 IEEE.

two agents having the same trait value for a given feature. In extensions of Axelrod's work, these changes with q and F are often characterised as a phase-transition; the number of cultures at stabilisation jumps abruptly from a single monoculture to cultural fragmentation around a critical point of q [35]. These abrupt changes in the overall nature of the stabilised simulation do not occur when using our model. The multiple sources of influence have a mediating effect on the number of cultural regions, allowing cultural diversity and mixing.

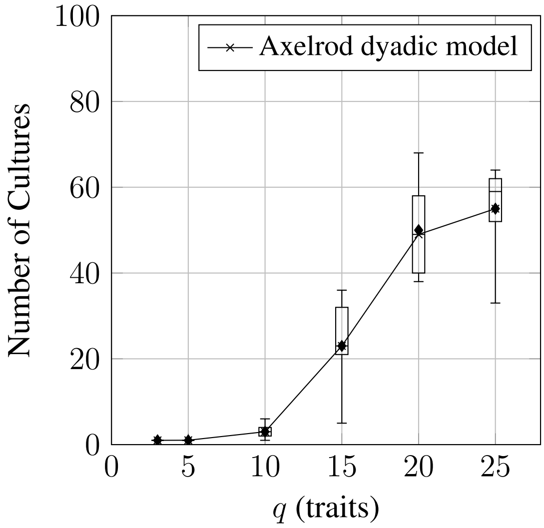

(a) Varying q. F = 5, Axelrod's model

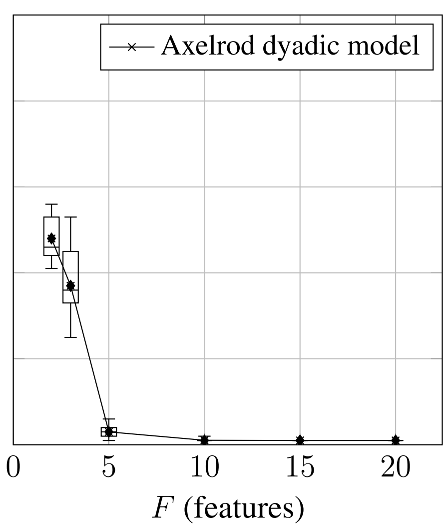

(b) Varying F. q = 10, Axelrod's model

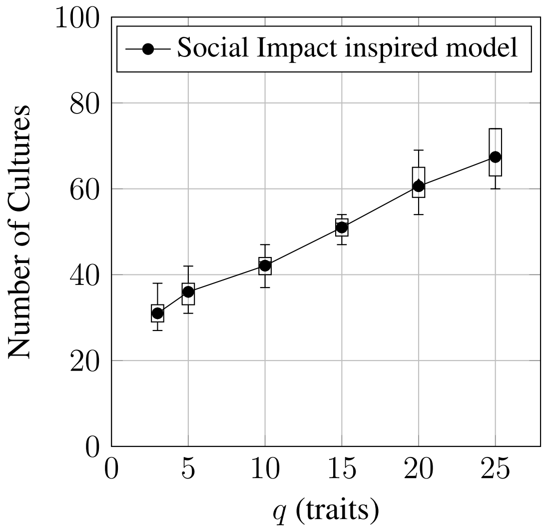

(c) Varying q. F = 5, our model

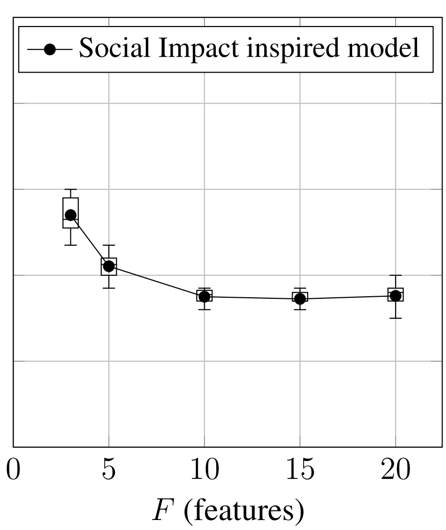

(d) Varying F. q = 10, our model

Figure 3.4: The effect of increasing the number of possible traits q on the number of cultural regions (3.4a, 3.4c), and the effect of increasing the number of features F on the number of cultural regions (3.4b, 3.4d) for both Axelrod's dyadic model of cultural influence and the Social Impact inspired model of peer pressure. Complete monoculture or complete cultural fragmentation are rare in our model; instead overlapping cultures may persist.

3.3.2 Polarisation, homogenisation, and the mixing of adjacent cultures

As mentioned in Section 3.3.1, typical measures for studies based on Axelrod's model are based on the number of cultures: agents with identical traits. This focus on region size is perhaps a consequence of the depiction of the simulation on a grid, similar to cellular automata. We however envisage our model as one of a network graph, and as such we are interested in both the state of the vertices and the edges which connect them.

Measures of cultural region number and size are often referred to as the degree of polarisation [6], however in the Axelrod model it is arguable that these metrics better illustrate a degree of fracture within an *always* globally polarised population. In such a dyadic model all agents are always either completely similar or dissimilar to their neighbours at stabilisation; cultural regions and their sizes are neatly delineated. In a model of compound influence such as ours, these metrics do not fully convey the dynamics taking place. At stabilisation, agents often hold traits from multiple neighbouring cultures simultaneously and therefore cultures counted as distinct may in fact have many traits in common. Similarly, the largest region size rarely adequately indicates the amount of fracture within groups, or the amount of cultural overlap between the largest region and others. A few dominant traits may persist across several otherwise different cultures.

Consequently, we introduce measures that record the state of the edge between nodes:

- **Homogenised** - an edge between two agents with identical traits, i.e. $similarity_{i,j} = 1$
- **Polarised** - an edge between two agents with zero traits in common, i.e. $similarity_{i,j} = 0$

- **Mixed** - an edge between two agents with some, but not all, traits in common, i.e. $1 > similarity_{i,j} > 0$

The total number of each of these edge types is normalised by the number of total nodes. Together these metrics give an indication of the amount of cultural polarisation, or on the other hand, the degree of cultural pluralism, present in the network. Where the number of 'mixed' edges is low (or zero), the simulation has ended in a highly polarised state: most neighbours are either identical or completely polarised. A high degree of 'homogenised' edges suggests a small number of large cultures, or a single monoculture.

Figure 3.5 gives an illustration of these agent populations represented as network graphs rather than sites in a grid. In this image, the difference in edge status gives us a flavour of the different characteristics of our model next to Axelrod's. Where in the dyadic model polarised edges form the boundaries between distinct cultures, edge statuses in our model are more varied and depend on local context. Two neighbours may share nothing in common, but may both have traits in common with another neighbour - a scenario not possible in Axelrod's model. Some of these 'local contexts' are examined in more detail in Section 4.2.

The proportions of polarised, homogenised, and mixed links within our model are not so closely tied to the overall number of unique cultures (Figure 3.6). While some different parameter sets may result in a very similar number of cultures, the degrees of cultural mixing or polarisation can be quite different. When increasing F, essentially the number of things on which it is possible for agents to have in common, a single monoculture is a foregone conclusion in Axelrod's model as homophily drives convergence (Figure 3.6c). Less so in our model, where it can instead drive diversity (Figure 3.6d).

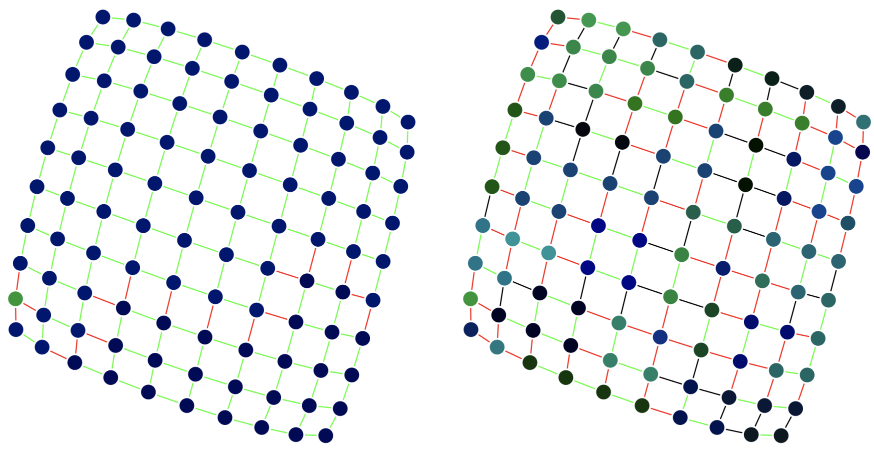

(a) Axelrod's model (b) Our model

Figure 3.5: Network graph representations of Axelrod's model versus ours in grid lattices. Axelrod's model (3.5a) stabilises with a small number of cultural regions, each containing identical agents. Edges between these identical nodes are considered 'homogenised', and depicted in green. Borders between cultural regions are comprised of polarised edges (red); the cultures can share nothing in common or one would assimilate the other. Mixed links (black) between agents with something, but not everything, in common are not possible in Axelrod's model at stabilisation. However, in our model (3.5b) it is possible for the simulation to stabilise with agents sharing a mix of traits from different neighbours. In this representation, vertex colours have been created simply using RGB values drawn from the trait values held. Nodes of different cultures will have slightly different colours.

3.3.3 Time to stabilisation

Direct comparison of time-to-stability between our model and Axelrod's is difficult as the typical number of interactions are on very different scales. A single interaction in our model entails an agent examining all of its neighbours and their features simultaneously, and potentially copying multiple features in one go. The dyadic model on the other hand requires several interactions before it can be said to have been influenced by

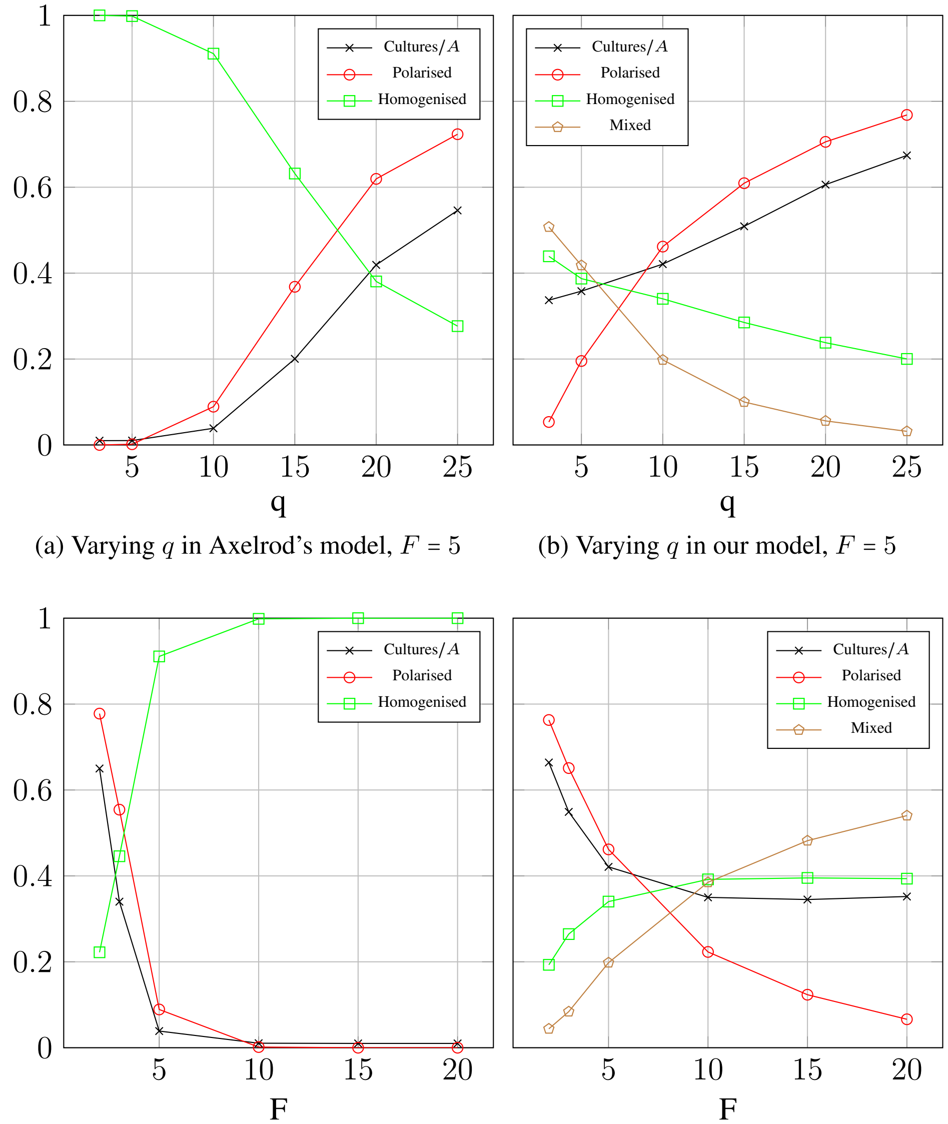

(a) Varying q in Axelrod's model, $F = 5$ (b) Varying q in our model, $F = 5$

(c) Varying F in Axelrod's model, $q = 10$ (d) Varying F in our model, $q = 10$

Figure 3.6: Edge statuses in our model versus Axelrod's on grid lattices. The usual increases in cultural fracture with increased q, and cultural convergence with increased F, are retained in our model, but with a mediated effect. The influence of multiple neighbours simultaneously prevents the simulation from reaching a state of complete monoculture or complete cultural fracture. The abrupt changes in state reminiscent of phase transitions (3.6c) are absent in our model (3.6d). Where polarisation and number of cultures are closely related in results from Axelrod's model, in simulations of our model the statuses of edges may change even where the total number of distinct cultures does not (3.6d).

all its neighbours, as it only interacts with one neighbour on one feature per iteration.

As a result our model does appear to have a more consistent run time due to the stabilising effects of simultaneous compound influence. Across the parameter sets illustrated in Figure 3.4, our model typically stabilises at between 550 and 1800 iterations, with a relative standard deviation of 22%. Axelrod's model typically stabilises at between 1800 and 225500 iterations with a relative standard deviation of 72%.

Descriptions of how our model, and Axelrod's model, stabilise are given in sections 3.2.1 and 2.3.1 respectively.

3.3.4 Comparison with Flache and Macy's social influence model

For comparison, we also run an adaptation of Flache and Macy's model [59] with homophily and social influence activated, but noise and selection error removed. Under these parameters, this is also an Axelrod-based culture model of multiple influences, but these are calculated in quite different ways.

In replicating an adaptation of Flache and Macy's model we noticed other dynamics which differ from ours, ones which could be problematic in the extension of a model to consider status, directed networks, and distance.

The paper [59] suggests that with zero noise, *'equilibrium is guaranteed and easily tested (all pairs of agents are either identical or dissimilar to one another on all features; hence, no further change is possible and iteration can be terminated).'* This suggests a similar stop condition to Axelrod's model. However, in the specific parameters we chose when reproducing the model (homophily, social influence, zero noise and zero selection error), we find that this state of complete polarisation is not always obtained. The simulation may persist in the same state for many millions of iterations, while its agents sometimes still share some cultural overlap with neighbours of a different culture. Upon examination, it became apparent that there were certain local conditions under which cultures could overlap. Where agents have at least some

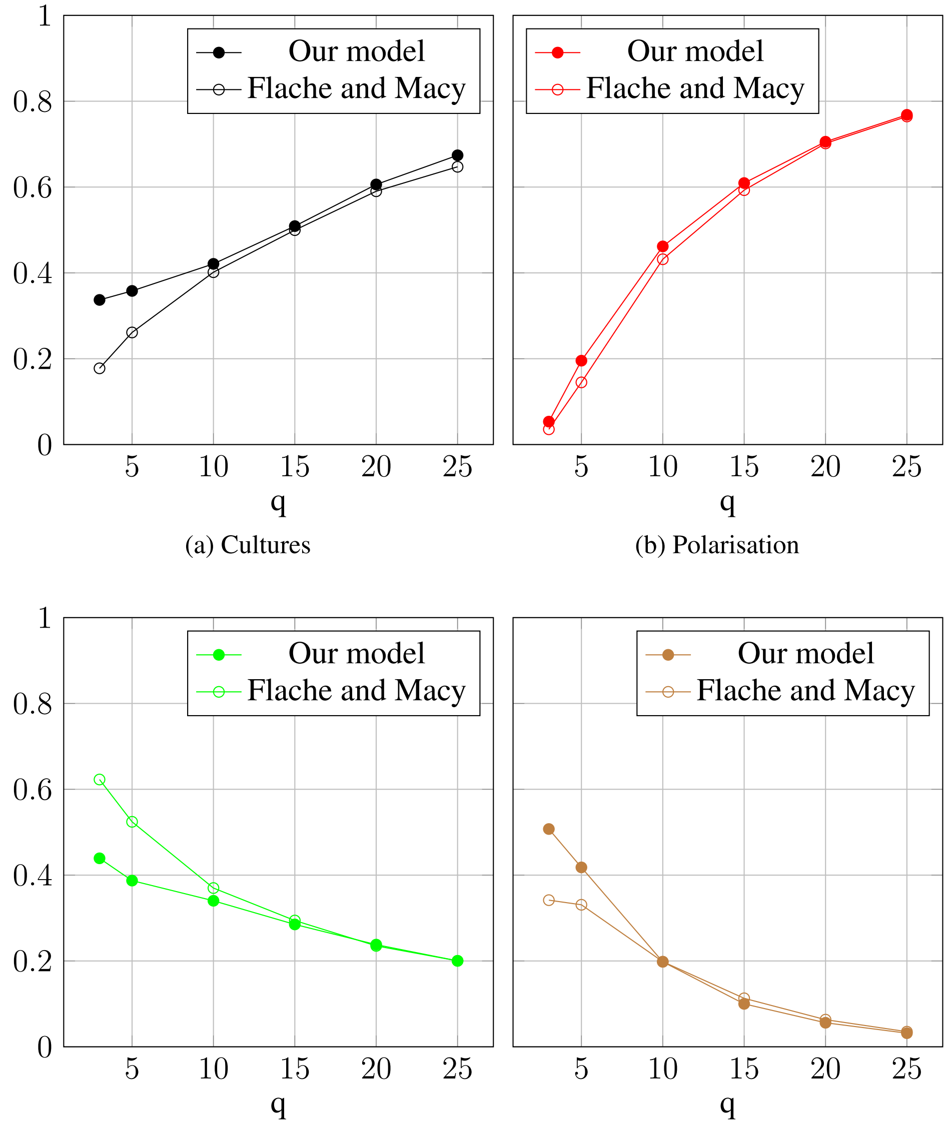

(a) Cultures (b) Polarisation

(c) Homogenisation (d) Mixing

Figure 3.7: Comparison of edge statuses (see Section 3.3.2) between our model and Flache and Macy's. Both are models of multiple social influence, and both show similar averaged results. However, our model shows greater cultural mixing and less convergence, particularly when the number of traits q is lower.

neighbours with which they share an identical cultural identity, they may also have other neighbours with which they share some traits - but from whom social influence will never be enough to counter that of the identical neighbours.

Results of this adapted Flache and Macy model bear more similarity to our model than Axelrod's, confirming the significance of multiple social influences (Figure 3.7). However, there are subtle differences. Our model displays greater cultural mixing (Figure 3.7d) and less convergence (Figures 3.7a, 3.7c) than Flache and Macy's, particularly at low values of q.

A possible contributing reason is one remarked upon above; for mixed links to exist in the Flache and Macy model, each of the agents linked must also have neighbours to which they are completely similar. It is not possible for an agent to only have mixed links without also having some homogenised. The simulation forms neat cultural regions of homogeneity, separated by borders of either polarisation or cultural mixing (Figure 3.8a). This is not the case for our model. It is perfectly possible for an agent to stabilise while having mixed links, and no homogenised links. They can take on a mix of the cultures around them without a requirement of being tied to another identical agent (Figure 3.8b).

These differing dynamics make it less clear as to when this adaptation of Flache and Macy's model can be said to have stabilised, and an algorithmic detection of equilibrium in the simulation is less straightforward than either our model or Axelrod's. In fact, in Section 7.4.2 we find instances where the Flache-Macy model does not seem to reach a definite equilibrium within a reasonable time frame. In our adaptation we included a stop-condition where if each and every agent has had an opportunity to copy x times and does not do so, the simulation can be said to have stabilised. x can be tuned, and we found no traits were copied in the interval $1000 < x < 10000$.

We believe that our formulation of multiple influences is a more explicit modelling of *number* from Social Impact Theory. The simultaneous consideration of all neighbours within a network distance makes it easier for us to use our model to isolate the different factors of Social Impact Theory. Also, it should be more straightforward to extend our model to consider factors such as weighted, directed edges (see Chapter 5) and distance along a network path (Chapter 6).

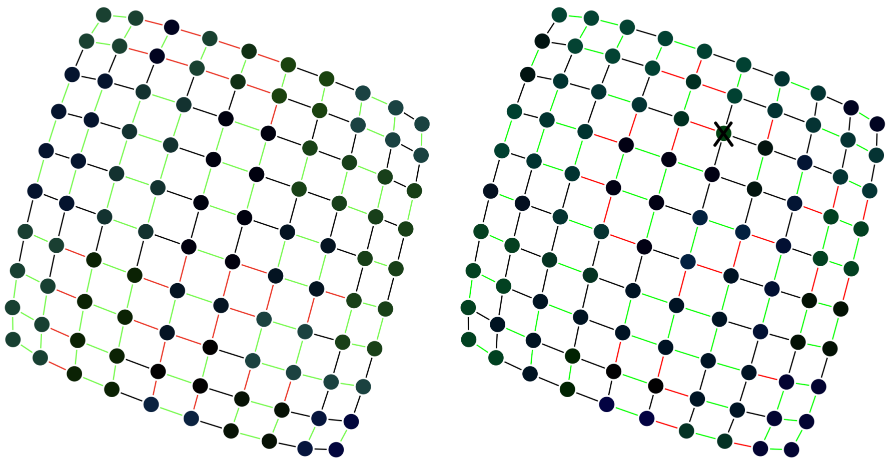

(a) Flache & Macy model (b) Our model of multiple influences

Figure 3.8: Network representations of Flache and Macy's model (left) versus our model (right) when run on the same starting traits. In the Flache and Macy simulation, at stabilisation mixed links can only form the borders of cultural regions where agents are identical. Where a node has a mixed edge (black), it must also have at least one other edge which is homogenised (green). In our model, agents are free to hold overlapping traits with several cultures at once, and do not need the reinforcement of an identical neighbour. An example is given by the node marked with an X (3.8b), which has mixed and polarised edges to its neighbours but is not identical to any of them.

3.4 Conclusions

In this chapter, we extended Axelrod's culture model by embedding Latané's Social Impact Theory as the basis for social influence between agents. In particular, in this chapter we have focused on the addition of the *number* social force from Social Impact Theory, in a mechanism that simulates 'peer pressure': multiple simultaneous influences upon an individual.

The most significant outcome of our extension is the emergence and maintenance of cultural plurality, rather than inevitable polarisation. The cultural zones identified in

Axelrod's work do not become homogeneous, instead retaining a number of different cultures within with overlapping features. Arguably, this phenomenon is frequently observed in the real human social world. Adjacent cultures are often made up of various features - some in common, some distinct - without necessarily converging to one monoculture or diverging to a state of polarisation. Minority languages persist despite pressure from neighbours, even those with which much is held in common. The local dynamics allowing this state of equilibrium are a form of social reinforcement, where an agent may retain its traits in the face of pressure from others because of the competing influence of other neighbours who hold a portion of the existing traits.

The use of the social force of *number* from Social Impact Theory offers a candidate explanation - social reinforcement from multiple similar sources - for the persistence of different overlapping cultures in the global population. Note however, that mere exposure to another culture is not sufficient to allow mixing; agents need something in common. There are observed instances where exposure to an opposing view does *not* encourage an individual to move their own opinions in that direction [8, 10]. Also, exposure to a different but compatible culture without any social reinforcement of existing traits would result in the agent changing to adopt the new culture in its entirety. There may be analogies here with isolated individuals being more prone to radicalisation when brought into contact with extremists of similar views.

It should be noted that we did not seek to create a model that allowed or encouraged cultural mixing and diversity; our aim was simply to test whether such an extension of Axelrod's model would produce substantially different results in order to answer research question RQ1. These results indicate that applying Social Impact Theory in such a way can indeed cause considerable differences in the global state of the system.

While we acknowledge that there are likely alternative approaches to modelling the combined effects of peer influence, our approach offers a useful benchmark and basis for future work. Its clear criteria for stabilisation, and separation of aspects of *strength*, *immediacy* and *number* allow for extension in each of these areas. A purely deterministic model was considered, one in which all agents would consider and update their traits simultaneously. Such a simulation would, under certain specific conditions, fall into a state of oscillation, similar in some ways to various cellular automata. However, we do not know if such a deterministic system would be particularly realistic[2] in the application of Social Impact Theory to culture models.

The persistence of cultural plurality in our simulation appears to be a consequence of agents becoming 'stuck' between two cultures, and adopting traits from either. Social reinforcement from different contacts allows an individual to retain cultural or behavioural attributes even where other contacts may act differently. A potential explanation for this dynamic lies in the theory of Complex Contagions [38], the idea that behaviours do not spread as easily as contagious diseases but instead require multiple sources of exposure for transmission. In Chapter 4 we explore this idea, as well as broader implications of network structure, in order to gain a greater understanding of the local dynamics that are causing this macroscopic phenomenon.

[2]Despite this thesis being submitted as part of the fulfilment of the qualification Doctor of Philosophy, I feel a discussion on the existence of free will is probably beyond the scope of this thesis.

Chapter 4

Cultural Diffusion as Network Contagion

In this chapter, we compare the behaviours of both our multiple-influence model introduced in Chapter 3 and Axelrod's to studies of contagion dynamics, principally those that focus on the difference between *simple* and *complex* contagions [38]. As our model is one of multiple simultaneous influences, it may show similar differences to Axelrod's model of dyadic influence as models of complex contagions show to those of simple contagions. By comparing our work to contagion studies on small-world networks, we aim to test network properties that have particular effects on our model. The use of rewired ring lattices allows us to control and examine effects of average path length, local degree and clustering coefficient. This chapter supports research questions RQ2-RQ4.

4.1 Simple and Complex Contagions

Robert Axelrod's model [6] is one of cultural spread and polarisation within a cellular-automata-like grid. While his work primarily concerns the macroscopic polarisation and segregation that can occur from micro-level rules, it nevertheless relies on the basic diffusion of traits throughout a network - even if it is initially a very simple lattice. Many subsequent extensions, including ours, use techniques from network science and graph theory in their analysis [37, 85, 143]. The copying of traits within culture models also bears some similarity to models of *contagion*.

In Section 2.2.3 we gave a brief overview of contagion models in general, and noted the many studies which apply network science to the problem. Granovetter's influential work on 'weak ties' stressed the importance of long, casual connections acting as bridges between communities [69]. Watts and Strogatz created 'small world' networks by using just a small number of these bridging ties to act as shortcuts across a network, giving a low average path-length but retaining a high degree of clustering [160]. These structures support the spread of 'viral' contagions, by allowing short paths for a contagion to spread quickly throughout the network. However, it is disputed that social phenomena such as behaviours and norms spread as easily as pathogens; there are many cases where spread of a trait only occurs when there are multiple sources of influence (see Section 2.2.3). Where the threshold to adoption is higher, Damon Centola characterises this as a *'complex contagion'* [38, 36].

The definitive difference between a *simple* and *complex* contagion is the minimum number of contacts required for transmission. We cover this theory in more detail in Section 2.2.3. In simple contagions, such as infectious diseases, a single interaction with a single contact may be sufficient for a trait to spread. A complex contagion on the other hand, requires multiple sources for transmission to occur. Centola argues that this is often the case for the adoption of new behaviours, innovations, and social norms [36]. Where contagions are complex this can impact widely held assumptions regarding how behaviours spread 'virally' across the bridging ties of a small-world network.

4.1.1 How the diffusion of simple and complex contagions differ

In [38], Centola and Macy illustrate some of these differences in how a complex contagion may spread through a network in comparison to a simple contagion. They apply the Watts-Strogatz model [160] to complex contagions, investigating whether the small-world principle generalises from the spread of information or disease to the diffusion of collective behaviour. Nodes are given a simple status: activated or unactivated. Nodes change from unactivated to activated when the number of their activated neighbours exceeds a threshold τ. Activated nodes never return to being unactivated.

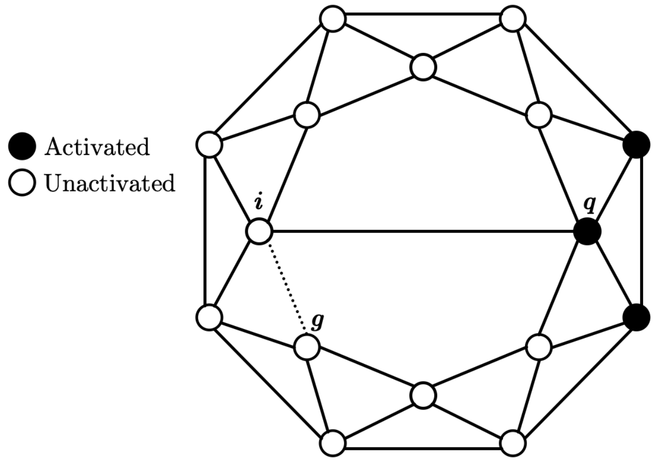

Figure 4.1: Simple contagion in a Watts-Strogatz graph: The edge between gi has been rewired to iq, allowing a simple contagion from q to i; i will become activated through contact with q. However, such a narrow bridge across the ring would not allow even a minimally complex contagion of $\tau = 2$.

In threshold models the threshold may often be expressed either as a whole number (of neighbours), or as a fraction $\tau = a/z$, where a is the number of activated neighbours required for transmission and z is the total number of neighbours. For example, when expressed as a number $\tau = 3$, at least three sources are needed for activation - regardless of the number of other neighbours who may be unactivated. When expressed as a fraction, for example $\tau = 3/12$, three out of twelve neighbours must be activated for the node to be activated itself. Where complex contagions and fractional thresholds are involved, Centola and Macy find there is often a qualitative difference between $a = 1$ and $a > 1$, even when proportions are identical. E.g. the fractional thresholds $1/4$ and $4/16$ may display different behaviours: $1/4$ is a simple contagion, $4/16$ is a complex contagion. The numerators are important: where they are greater than 1, contagion is *complex*, requiring multiple sources.

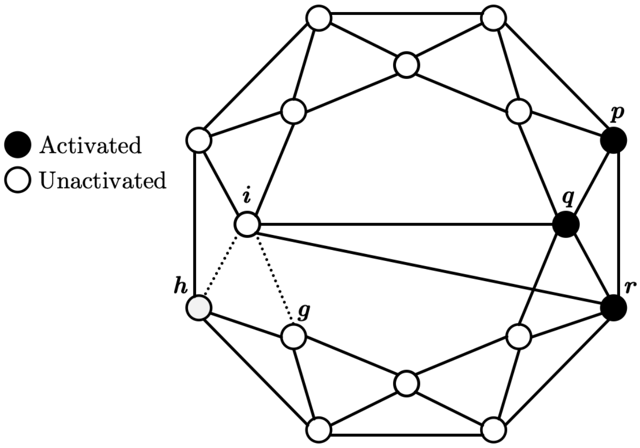

Figure 4.2: Complex contagion in a Watts-Strogatz graph: While the graph in Figure 4.1 would not allow complex contagion across the bridging tie, here an additional edge hi has been rewired to ir. Now a complex contagion of $\tau = 2$ (or fractionally, $\tau = 2/4$) may cross the wider bridge, and activate i. Note however, the growing erosion of the ring structure. Once i has been activated by q and r, it will not be able to spread the contagion to h or g as those links have been rewired. Node h will only be activated by a complex contagion spreading clockwise around the ring; its ties to those nodes on its other side have deteriorated below the level needed to transmit a complex contagion of $\tau = 2$.

Centola and Macy state that for complex contagions, the bridges between neighbourhoods across a graph must exceed a critical width, W_c, for contagion to occur. On simple ring lattices of degree 4, this is calculated $W_c = a(a+1)/2$. A seemingly obvious way of increasing bridge widths across the ring is to increase the amount of random rewiring (as Watts and Strogatz did in their ring lattices [160]). However, as edges are randomly created across the ring they are often removed from the ring itself, eventually lowering the overall clustering of the network and hindering the spread of complex contagions locally. Where p is the probability of rewiring a link, they find that increasing p causes a U-shaped effect on propagation of complex contagions. Increasing randomisation of links will typically aid propagation of a simple contagion, even with only a small increase in p. For complex contagions, a larger increase in p is needed to affect diffusion. As p rises the propagation time first drops, then rises (see [38], Figures 3 and 4). There is a 'small-world' window that may facilitate complex contagions, where path length is short but clustering remains high.

There are aspects of this theory of complex contagions that seem useful for explaining some behaviours of our model of multiple influences, and why diffusion of cultural traits is typically lower than that in Axelrod's. However, there are important differences between this simulation of complex contagions on ring lattices and our model. In Centola and Macy's initial investigation, every node has equal influence and an identical threshold. In our model, the similarity of nodes determines their influence, and hence the thresholds of their neighbours. While in Centola and Macy's complex contagion study, nodes have a simple binary status of activated or not; in our model agents typically have a broader set of features and possible values (traits) for each feature. In their model, nodes' statuses may only change in one direction: from unactivated to activated, and do so deterministically immediately after a threshold is exceeded. In our model traits can spread in either direction across an edge, and agents may adopt different traits - or readopt previously held traits - several times over the course of the simulation. Additionally, while some aspects of our copying mechanic are deterministic, agents are activated randomly. A node's neighbours may, by chance, change state several times before the node is called upon to copy traits themselves. In the next section, we give examples where our model of multiple influences may exhibit behaviours similar to both simple *and* complex contagions.

4.2 Our multiple-influence culture model exhibits both simple and complex contagious local behaviours

If comparing the diffusion of traits in culture models to the spreading dynamics of contagion models, it would be prudent to determine whether the copying mechanisms of those models more closely match single or complex contagions. Most studies of complex contagions [38] set a threshold for contagion - the number of activated neighbours needed for a susceptible node to adopt a trait. As we have seen (Section 4.1.1), for simple contagions this is always exactly one neighbour; for complex contagions either an absolute figure greater than one or a fraction of a node's neighbours where the numerator is greater than one. This threshold may be heterogeneous or homogeneous across all nodes but it is usually known and fixed from the start of the simulation. In the culture models we examine however, the state of each agent is frequently changing, and with it the agent's relationship with its neighbours. Whether an agent can copy from a neighbour is dependent on the traits both agents hold at that precise time-step.

Simple contagions in the Axelrod model

Nevertheless, it is clear that Axelrod's model is essentially one of simple contagion. While the state of agents affects the probability of copying, all interactions are dyadic. When activated, an agent only takes into account one random neighbour at that time (Figure 4.3). Whether the agent adopts a trait from that neighbour is a matter of chance, but the states of any other nodes are irrelevant to the interaction. There is no social reinforcement, no peer-pressure; the minimum number of sources required for adoption of a trait is one.

Simple and Complex contagions in our Multiple-influence culture model

In models of multiple simultaneous influences such as ours, thresholds for contagion vary and are dependent on the traits held by the agent and every other node that has influence over it. In a situation where an agent has a high degree of similarity with its neighbours, the threshold to change will be quite high (Figure 4.4). If one imagines a scenario where an individual has a great deal of affinity with their social group, any interloper would find it difficult to persuade that individual to change deep-seated

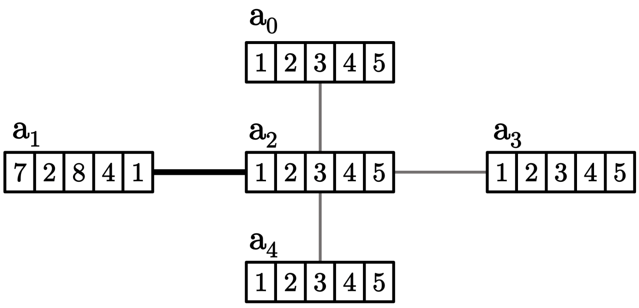

Figure 4.3: Simple contagion in the Axelrod model: a_2 is interacting with a_1. The similarity of a_2's other neighbours is irrelevant; a_2 is free to copy from a_1 as a single source.

behaviours, traits, or opinions away from those of its peers. The individual would be far more likely to adopt a new trait - perhaps a new fashion - if it first saw such a change in several others they deemed to be similar in values to themselves (Figure 4.5). This would be an example of homophily from multiple sources raising the minimum threshold to change, and thus making contagion complex in nature. This is a key dynamic in our model of multiple influences, and goes some way to explaining the differences in cultural dissemination between ours and Axelrod's model.

However, our model is not just one of complex contagion. There are circumstances where the threshold to change is low. In contrast to the scenario where an agent is surrounded by similar peers, an agent who has little similarity to its neighbours is more susceptible to copying traits from the few sources it identifies with (Figure 4.6). An analogy may be a person who feels they have little in common with their close network, those links existing instead perhaps for reasons of family, employment, or geographical proximity. Such a person may be more likely to be influenced by a contact it is not so alienated by, one with which it shares at least some common ground. Alternatively, they may become impressionable to distant but seemingly like-minded persons met online, who share similar interests; analogous to Granovetter's 'long-ties', or the rewired bridging links that allow simple contagions in a Watts-Strogatz graph. This could have implications for the modelling of issues such as loneliness, or even the radicalisation

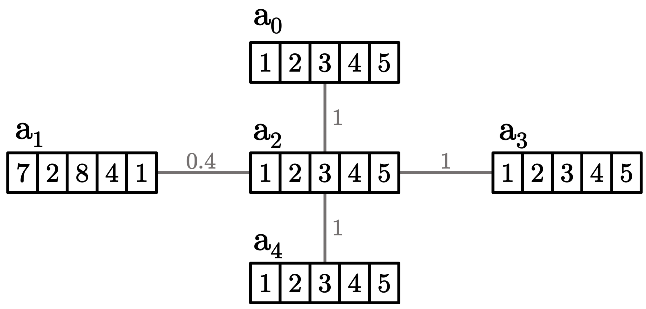

Figure 4.4: Social reinforcement in our model: It is not currently possible for a_1 to influence a_2; the influence of a_2's other neighbours is too great. In fact, a_2 will never change in this situation, unless its influential neighbours change first or it develops a substantial number of new links to other nodes. When considering feature 1, a_2's neighbours have either trait 1 or trait 7. The sum influence of the three neighbours with trait 1 is 3, whereas only a_1 has trait 7 and is only able to exert an influence of 0.4 on a_2. Thus the threshold for contagion is $\tau > 3/z$, requiring a complex contagion for change to occur.

of alienated persons.

These examples illustrate that our model is one of both complex *and* simple contagions. The threshold to change for any individual node is wholly dependent on local conditions. Areas of a graph with strongly connected nodes of a single monoculture may be almost (but not entirely) impossible for outsiders to make inroads into. Elsewhere in the same graph, locally dissimilar agents may be easily changed by others. Much of the population may sit in-between, adopting different elements of bordering cultures. A person can adopt different traits from their various types of contact - their colleagues, family, friendship groups, religious groups and so on. One may support a different football team to one's colleagues, have different hobbies and interests to their family, prefer different foods to their friends. As we have seen in Chapter 3, the cultural mixing which occurs in our model makes it distinctive from the Axelrod model on which it is based.

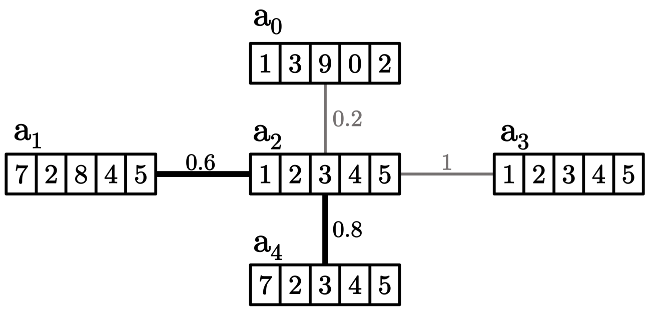

Figure 4.5: Complex contagion in our model of multiple influences: Even when an agent has a high degree of similarity to a neighbour, others can effect change. Here, agent a_2 is identical to its neighbour a_3. However, the combined influence of a_1 and a_4 is greater still. When considering feature 1, the trait score of trait 7 is 1.4 (from a_1 and a_4) versus a trait score of 1.2 for trait 1 (from a_0 and a_3. Therefore, at the next opportunity agent a_2 will adopt trait 7 as its first feature. This is an example of a successful complex contagion; a single neighbour would not be able to exert enough influence on a_2 to enact change, but multiple neighbours can.

4.3 Simulations of our multiple-influence culture model on ring lattices and small-world networks

In this section, we compare the macroscopic behaviours of our multiple-influence model and the dyadic Axelrod culture model on ring lattices and Watts-Strogatz networks. There are a number of reasons for choosing these network structures for analysis. Ring lattices allow us to increase the degree of all nodes in a heterogeneous fashion, so no one node has different connectivity characteristics to any other. Also, some network distance (i.e. graph diameter and average path length) is maintained. While these characteristics are shortened as a ring lattice increases its degree, the shortening effect is less than many generated random graphs.

Conversely, the rewiring of a ring lattice that occurs in a Watts-Strogatz graph allows us to rapidly shrink the diameter and average path length of the network while maintaining the same average degree, and a close-to-uniform degree distribution.

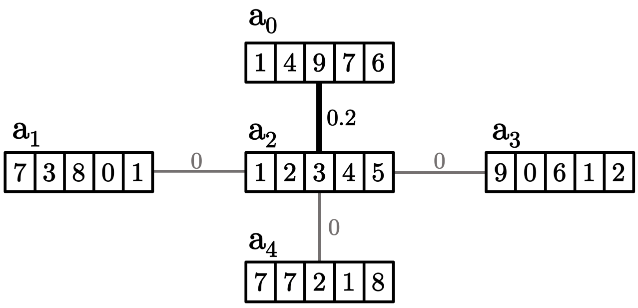

Figure 4.6: Simple contagion in our model of multiple influences: Despite having little in common, agent a_2 is influenced by agent a_0. Because a_2 is polarised with its other neighbours, its threshold to change is low, and when activated it will copy traits from a_0. In this situation, a single contact with any similarity at all can have an influence - an example of simple contagion.

In addition, these network structures have been used in other studies to compare simple vs complex contagions, most notably [38]. We expected that Axelrod's model would display behaviour more similar to that of a simple contagion model, and that our multiple-influence model would show some characteristics of complex contagion.

4.3.1 Methodology

We arranged 1024 agents in ring lattices of varying degrees: 4, 8, 12, and 24. Starting traits were assigned randomly to each node, where $F = 5, q = 10$. Both Axelrod's model and our model of multiple-influence were run on these generated structures. Runs were limited to 10 each per parameter set, due to limited resources. A single run of the simulation can often take several hours to stabilise. Despite the low number of runs, result metrics obtained showed a narrow distribution of values, low variance, a small range between minimum and maximum values, and similar median values to mean (see Figures 4.7, 4.8). Additionally, select instances of 20 runs returned very similar results to 10 runs on the same parameters. This gave us confidence that the results obtained were indicative of the general behaviours of the models. The same random seeds were used across different parameters so that each model had the same starting conditions.

We also rewired the above ring lattices using the Watts-Strogatz model, with probabilities p ranging from 0.001 to 1 that an edge would be rewired, where $p = 1$ is essentially a random graph. These graphs were generated using the Python library NetworkX's [139] connected_watts_strogatz_graph generator [121]; our aim being to maintain a single connected graph in each case with constant size and density, but increasing randomisation.

It should also be noted that when we talk of the *clustering co-efficient* of a graph, we use that calculated by NetworkX [120], based on the number of possible triangles through a node that exist. Although slightly different to the measure used in [160], results for ring lattices and small world networks are similar.

4.3.2 Results

Axelrod's model in ring lattices

As with grid lattices [6], there is a clear convergence of cultures as degree increases (Figure 4.7). In most structures, in addition to increasing density, increasing degree will often increase clustering and decrease average path length. The increased number of pathways for a trait to spread, in addition to the likely shortening of these paths, makes the convergence to a monoculture an expected behaviour in a model where traits are transmitted so easily.

Our multiple-influence culture model in ring lattices

In our model of multiple social influences, the number of distinct cultures in the population is also reduced as degree increases. However, there appears to be barely any effect

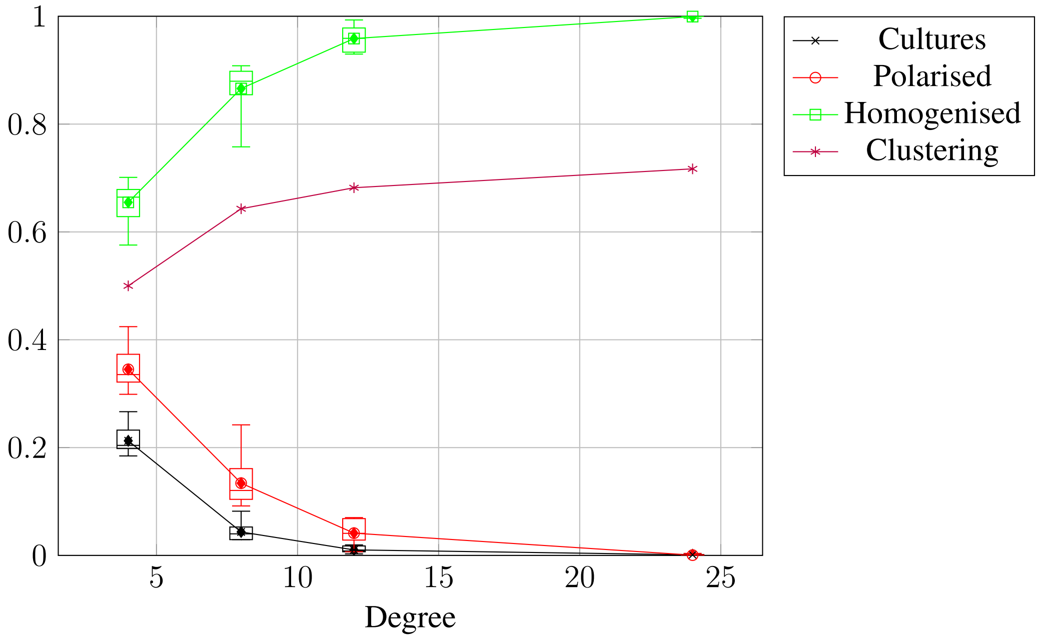

Figure 4.7: Homogenisation of Axelrod model as degree increases: As the node degree is uniformly increased in ring lattices, the Axelrod model stabilises at fewer cultures. Correspondingly, the amount of homogenised edges increase, and the amount of polarised edges decrease. Note: number of cultures has been normalised by number of nodes.

on the proportion of homogenised, polarised, and mixed links within the network (Figure 4.8). At first glance this may seem a contradiction; but while the number of cultures is normalised by the number of nodes, the portion of polarised/homogenised/mixed links are proportionate to the number of edges. As the number of edges increases, the ultimate number of homogenised/polarised/mixed links also increase at a fairly linear rate, resulting in very similar proportions in highly connected rings to lower connected rings. This is largely due to the nature of ring lattices; clustering is high throughout and some distance across the graph is maintained. Any new links are created in similar areas to existing links: they will not be bridges across the ring - the 'weak links' of which Granovetter wrote. Where existing links in an area are between identical agents, additional links in the same area are likely to be also. While some cultures may

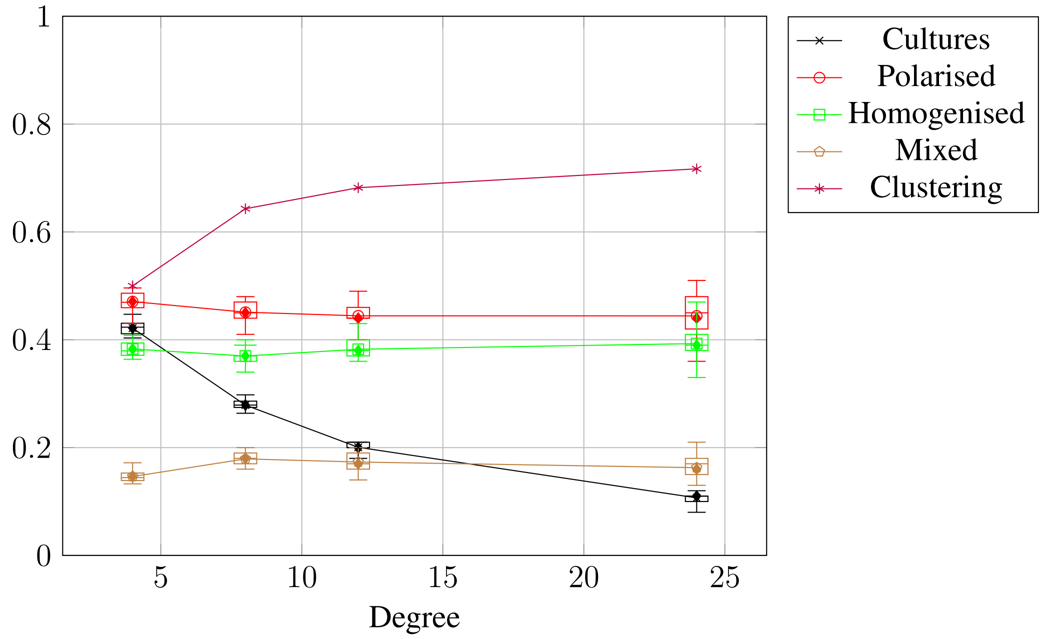

Figure 4.8: Cultural convergence in our multiple-influence model as degree increases: As average degree increases, clustering increases and cultures converge. However, the proportion of homogenised, polarised and mixed links remains largely the same. As more edges are added to increase degree, they become homogenised, polarised or mixed at roughly the same proportions as before.

converge, additional links are then created between that culture and other agents who share some, or no, traits in common. This maintains the percentage of homogenised, polarised, and mixed links in the overall network.

Axelrod's model in small-world networks

Where increasing the uniform degree of a ring lattice increases its clustering, rewiring even a small number of edges serves to greatly reduce average path lengths. As a model where diffusion of traits is analogous to simple contagion (see Section 4.2), any reduction in path length via 'short-cuts' across the ring will greatly increase cultural convergence (Figure 4.9). These links form bridges between cultures that may previously have been insulated from each other by other, non-compatible cultures; allowing traits to spread across the divide.

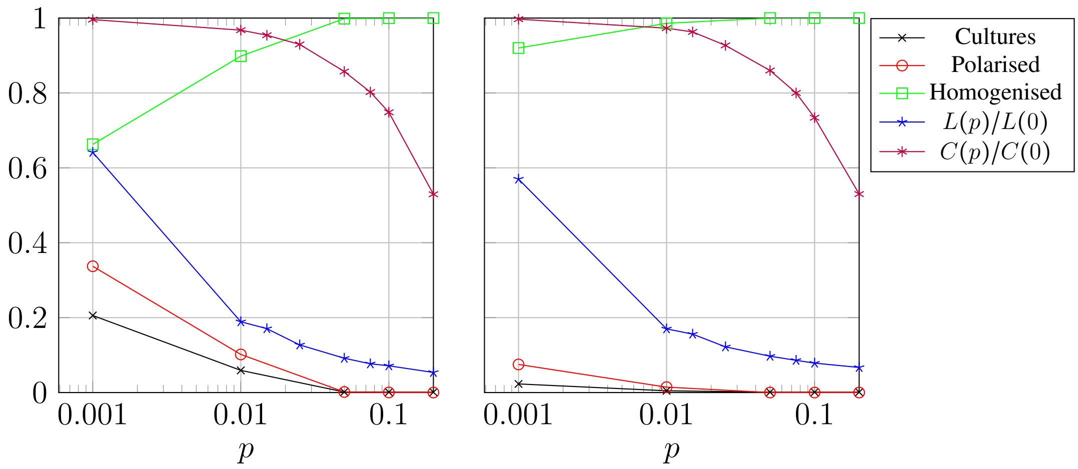

(a) Ring Degree 4 (b) Ring Degree 8

Figure 4.9: Homogenisation of Axelrod model as path length decreases: As the probability of rewiring p increases, the number of cultures quickly falls to a single monoculture. The ease at which agents copy from each other ensures that any shortcuts across the graph will greatly increase transmission of traits. When a graph is rewired enough, there are multiple pathways for traits to spread. Any cultures that might otherwise be isolated from other compatible cultures in another part of the network are far more likely to form connections, and thus assimilate each other into one culture. Note: the x-axis is a log scale of the probability p that an edge will be rewired. $L(p)/L(0)$ and $C(p)/C(0)$ are the average path length and clustering coefficient respectively, both normalised by their values when $p = 0$.

Our multiple-influence model in small-world networks

Increasing the uniform degree of a ring lattice showed that increasing clustering can have an effect on the cultural convergence of our model (Figure 4.8). As links are rewired across the ring, our model initially shows little change in homogenisation or cultural diversity even as average path length is reduced. It is only when clustering drops significantly that there are discernible changes in these measures. If our model contains behaviours similar to that of complex contagions (Section 4.2) then a relatively low sensitivity to changes in average path length is intuitive; single bridges across a ring are not usually sufficient to overcome the social reinforcement of an agent's often stronger (and more numerous) local ties.

It is not possible to directly compare our results to Centola and Macy's results [38] on convergence as p increases. As a model of contagion, their results concern the time-steps needed for a contagion to spread to the entire network. In our multi-dimensional feature model, traits or cultures rarely if ever spread to the entire network; instead we look to the differences between agents. However, a key result from their measuring of time-to-convergence was that increasing randomisation p may initially aid convergence, before ultimately impeding it when edge placement becomes more random and clustering falls (see [38], Figures 3 and 4). Were similar dynamics the sole cause of cultural diffusion in our model, we would expect to see a similarly U-shaped trend in either the number of cultures or degree of polarisation in our model.

Cultures do rise significantly when clustering drops, indicating decreased convergence at higher levels of p; a similar phenomenon to increased time-to-convergence in Centola and Macy's study. However, there is no discernible 'U-shape' or significant reduction in the number of cultures before they rise (Figure 4.10). There are instances of tiny fluctuations in the number of cultures between $p = 0.001$ and $p = 0.1$, but these are so small as to be likely caused by the random starting traits for each run. The 'small-world window' of high clustering and low path length does not seem to have much effect on our model.

Changes in polarisation as p increases are negligible, and likely not significant. For some higher degree rings there is a slight increase in polarisation when p increases from 0 to 1, but it is less than 0.05 in each case. Homogenisation on the other hand drops noticeably as the number of cultures rise. The explanation for homogenisation dropping and polarisation remaining largely steady, is that cultural mixing increases as clustered edges are rewired to nodes further away in the ring. It appears nodes are retaining some traits from their local neighbours, while also adopting others from

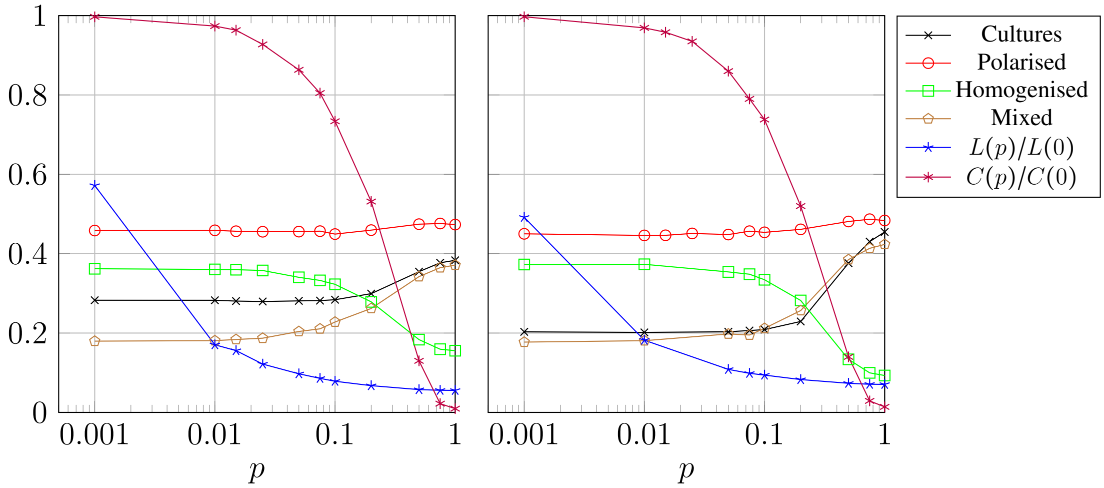

(a) Ring Degree 8 (b) Ring Degree 12

Figure 4.10: Increased cultural diversity of our model as clustering decreases: Unlike in some existing models of contagion, there is no increased cultural convergence in a 'small-world' network where clustering is high and average path length is low. When clustering falls to lower levels, cultural diversity and mixing increase. Note: the x-axis is a log scale of the probability p that an edge will be rewired. $L(p)/L(0)$ and $C(p)/C(0)$ are the average path length and clustering coefficient respectively, both normalised by their values when p = 0.

elsewhere in the ring lattice.

4.3.3 Discussion

Results from Axelrod's model on ring lattices and small-world networks show patterns in common with studies of simple contagions. Increasing either the average degree of the ring lattice, or the amount of random rewiring, causes greater cultural convergence and homogenisation. There are however differences, due to the multi-dimensional features of the Axelrod cultural model in contrast to the binary nature of many studies on network contagion. Transmission in Axelrod's model is also bi-directional; an agent may change its feature back and fore between traits before the simulation stabilises; in contagion models an activated node cannot usually 'catch' being deactivated. Whereas those studies typically converge to a single state of either activated or not, Axelrod's model often stabilises with several polarised cultures remaining. These cultures are amalgamated into one when graph diameter is so low that cultures are unlikely to be isolated from others they share a trait with (Figure 4.11).

While homophily is a key mechanic in Axelrod's model, and it influences which agents are more likely to spread their traits locally, if there is any similarity at all then eventually agents will tend to become the same. This is not the case for our model, where agents can find themselves between competing cultures, as we have shown in Chapter 3. There are some resulting behaviours analogous to complex contagions, such as the need for multiple sources of influence to effect change (see Section 4.2). As with comparing behaviours in Axelrod's model to simple contagions, multi-dimensional features of agents restrict the application of contagion dynamics. In our model this is even more pronounced, as agents may copy a variety of different traits from different neighbours and end in a 'mixed' state (Figure 4.12).

Additionally, there are cases using our model where traits spread very easily as if by simple contagion (Section 4.2). Even within the same neighbourhood, where there is a mix of traits for a feature an agent may change their trait very easily from a low number of influencers. Yet for other features it may require strong sources of influence to overcome existing social reinforcement.

These distinctions in local interactions cause differences in global network behaviour. While increasing rewiring in [38] had a nonmonotonic 'U-shaped' effect on propagation, there is no discernible increase in convergence in our model at any point. Our model shows an increase in cultural mixing and diversity as links are rewired toward a more random graph, and away from local clustering. This may have parallels in real life social networks. It naturally follows that if individuals are exposed to more cultural traits, cultural diversity may increase. Individuals may retain a large number of traits from their closest, clustered links - the 'strong ties' Granovetter described, often close

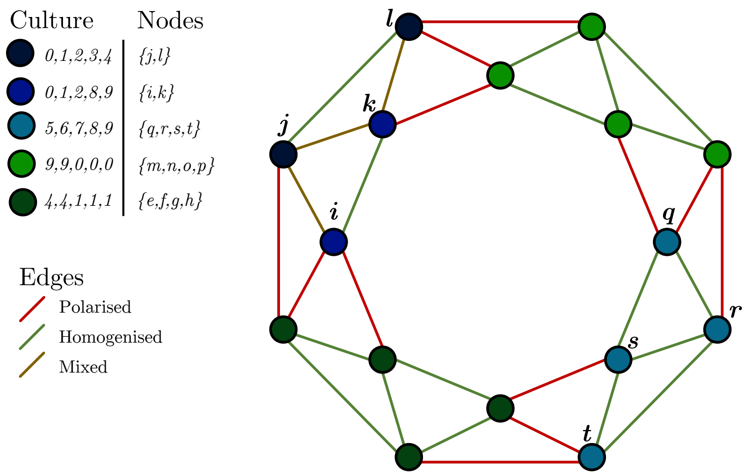

(a) A starting ring lattice with five cultures

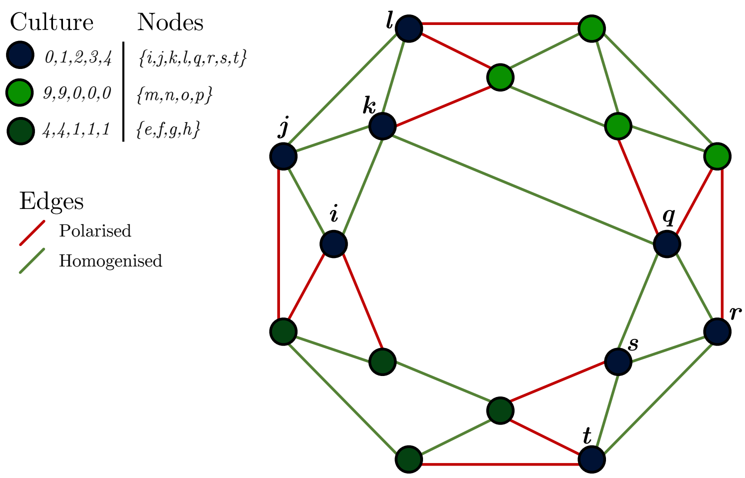

(b) The ring lattice in 4.11a, but with an edge rewired to kq, then run on Axelrod's model.

Figure 4.11: Dyadic copying as a simple contagion in a small Watts-Strogatz example graph: Should Axelrod's model be run on 4.11a, the connected cultures $\{i, k\}$ and $\{j, l\}$ will converge to one culture. In 4.11b this has occurred, and due to the edge kq has spread to the opposite side of the lattice. In any 'cultural region' of compatible overlapping cultures, the cultures will inevitably converge to one. Here, culture $0, 1, 2, 3, 4$ has prevailed, but any combination of traits existing in the cultural region $\{i, j, k, l, q, r, s, t\}$ may have possibly succeeded.

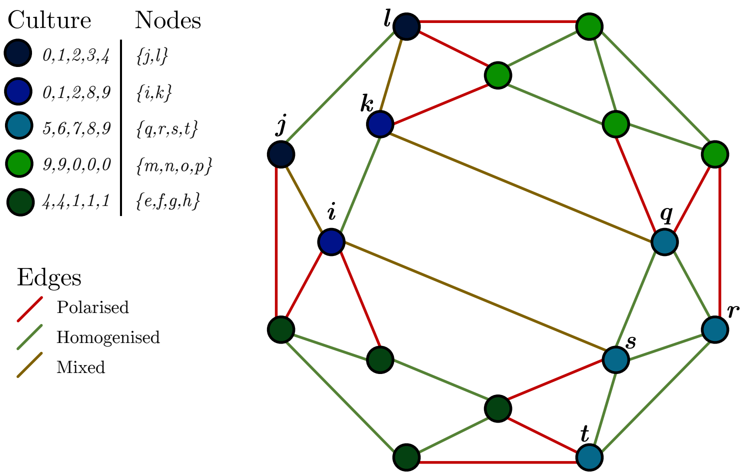

Figure 4.12: Cultural mixing in a small Watts-Strogatz example graph: Using our model of multiple influences, cultural diversity and mixing can persist even when adjacent nodes share traits in common. Here, nodes i and k retain a majority of traits common to their area of the ring lattice (j and k), yet also share other traits with the culture found at $\{q, r, s, t\}$. As rewiring increases, more mixed links are possible (see Figure 4.10).

family members and friends. However, as they broaden their horizons, they may adopt traits from other cultures they are exposed to (Figure 4.12). In our model, this tends to happen when clustering is lowered, which could occur either through removal of local clustered links or a greater proportion of ties being added to acquaintances across the network.

4.4 Conclusions

In this chapter we have examined the copying dynamics of our model within local neighbourhoods, and how they may compare to theories of simple and complex contagions. We then ran simulations on larger networks to examine the macroscopic effects these behaviours may have on cultural spread in a broader population.

In investigating research questions RQ2 and RQ3, we undertook neighbourhood-level analysis (Section 4.2) which showed that in our model of multiple influences, local conditions can allow copying behaviour analogous to either simple or complex contagions. In different circumstances traits may spread very easily, slowly, or not at all.

To answer research question RQ4, we undertook simulations in ring lattices and Watts-Strogatz ‘small-world’ networks (Section 4.3). These simulations show that the local dynamics described above manifest themselves as an increase in cultural diversity as links are formed across the network, but that these bridges usually need to be wide. This is in contrast to previous studies, which typically show an increase in convergence as edges are rewired and average path length falls. Our model appears to be more sensitive to changes in clustering than changes in average path length.

Our model may exhibit some characteristics similar to the concept of *interacting contagions* [99], in that the presence of other traits in a neighbour may increase propagation. However, it differs in that transmission is not dependent on a neighbour holding a particular trait, only that *any* traits be commonly held. Our findings may add weight to considerations that contagions may not be able to be neatly divided into simple and complex, and are dependent on local conditions. As Hébert-Dufresne et al observe, *"This mechanistic difference creates a false dichotomy, forcing us to choose the mechanism we think best describes the reality of a given contagion. In practice, the context of transmission events always matters."* [73]

The sensitivity to clustering in our model could have implications for the development of local ‘teams’ and shorter paths of influence in organisational structures, which we examine in Chapter 5.

Chapter 5

The effect of compound social influence in hierarchical structures

In this chapter, we extend our implementation of the social force of *strength* in Social Impact Theory to include hierarchical *status* in addition to the existing homophily. In particular, we examine this in pyramidical network structures as generalisations of organisational hierarchies. Where influence is from multiple sources, as in our model, will a higher status individual's influence be countered by the greater number of their subordinates?

The results of this chapter have been previously published in part in [113], and in a paper currently in preparation [114]. This chapter supports research question RQ5.

5.1 Organisational culture

Organisations that support similar functions or similar missions are well known to frequently exhibit very different cultures. Such culture can be an important determinant for an organisation's identity [72, 126] and its effective functioning [52]. As a term, *'organisational culture'* can have different connotations. It may be thought of as the values, norms and assumptions held by a group's members, regardless of whether they were intentionally set; alternatively, it is sometimes used to describe those values an organisation's leaders wish to instil and propagate [136]. As we have done throughout this thesis, we consider culture in a similarly broad sense as in Axelrod's work [6]; that is, anything over which individuals may influence each other, such as beliefs or behaviours. Other popular usages may equate organisational culture with *corporate* or *business* culture. We again use the term organisation broadly: as potentially any group of connected members who may influence each other. We are however principally concerned with those groups that contain some form of hierarchy.

5.2 Hierarchy and status

In alignment with the meaning of the word itself, all "organisations" have some form of structure for coordination, whether formal or informal, and hierarchy proves to be a prevalent and often defining feature of organisations in many operational contexts [97]. The notion of a hierarchy within an organisation may invoke an image of a formal, clearly defined hierarchy; however even where an explicit hierarchy appears to be absent, an informal hierarchy often develops [4, 53, 123].

There are numerous ways in which an individual in a hierarchy may be seen to be ranked higher than a subordinate. Building on French, Raven and Cartwright's work on social power [63], Magee and Galinsky make a distinction between power as the control of resources; and *status*, the amount of respect accorded by others [97]. Mowday [116] defines authority as legitimate power based on formal position, whereas 'power' and 'influence' are used as broader terms referring to a general ability to change the actions of others. In this chapter, we will use the term *hierarchical status* to encompass any means (intentional or otherwise) by which a superior may have a greater ability to influence a subordinate than vice-versa.

Hierarchies are typically pyramidical in shape; subordinates outnumbering superiors. We use tree network structures (specifically k-ary trees) as generalisations of these pyramidical hierarchies. We seek to model the balance between the greater influence of a higher node, and the combined influence of the greater number of its subordinates.

This greater influence due to status will be modelled by extending our model to use directed and weighted edges, a description of which will be given in Section 5.3.

5.2.1 Tall versus flat hierarchical structures

A recurring phenomenon in organisational culture is that of firms or organisations claiming to have moved away from the ‘traditional’ hierarchy. The terminology used varies, from the ‘Lattice Organisation’ at W. L. Gore & Associates in the 1970s [98], to more recent attention given to so-called ‘holacracies’ at companies such as Valve and Zappos [19, 61, 147, 154]. Claims of the elimination of hierarchy are likely overstated; as previously mentioned, the apparent absence of a formal or planned hierarchy is not necessarily the absence of hierarchy altogether [4, 53, 123]. Where the phraseology may vary with time and fashion, a consistent description is that of a ‘flatter’ organisation, with less intermediary layers of management.

Studies on such tall versus flat organisational structures have focused on performance or decision making capabilities, and largely in a qualitative fashion within the fields of sociology, psychology, or business studies [33, 70]. Performance of structured hierarchies in contrast to more democratic teams has been found to be dependent both on the nature of the task and the number of levels in the hierarchy [32]. Wider tree structures typically have a reduced path length for communication, but any superiors that do exist will necessarily have responsibility for a larger number of subordinates, potentially becoming bottlenecks. A taller, narrower structure on the other hand may insulate and distance higher nodes from lower.

In terms of computational social modelling and network analysis, Stocker et al [150] applied an opinion formation model to both taller and flatter hierarchies, finding that the latter cause more fluctuations from consensus than the taller trees. The spread of corruption through a network was simulated by Nekovee et al [118] using an epidemic-based model on flat-vs-tall structures. However we are not aware of previous work that has used these flat-vs-tall tree structures in the context of culture-based modelling. In this chapter we aim to rectify this by applying our model to such structures, in turn modelling the possible formation and spread of culture in an organisational structure.

5.2.2 Teams: local clustering within a hierarchy

Despite the presence of hierarchical status, influence is rarely entirely top-down. Individuals may employ a variety of tactics to exert influence upwardly or horizontally [82, 164]. Some may seek to increase their influence by combining with their peers in trade unions, teams, or coalitions [23] for example.

As organisations grow and labour is divided, different departments and structures may form, and each may have elements of its own sub-culture [136]. These sub-groups may subvert a consistent overarching organisational culture, but can also lead to innovation. At Intel in the 1990s, members of higher management were opposed to turning PCI chipsets into a business. A team was assembled that “flew under the radar”, without close scrutiny from higher management, to build the case for PCI chipsets as a viable business - which ultimately succeeded [29].

In this chapter we contribute a method for generating *team* sub-structures, where nodes with a common direct super-ordinate may form a connected ‘team’. These sub-structures will also increase the amount of local clustering, which as we have seen in Chapter 4 seems to have a greater effect on our model than path length. In existing literature [117], formal groups within a hierarchy are often referred to as *teams*, and informal groups as *coalitions*. The term coalition as used in this context is slightly different to that of political or military contexts, where it often refers to more formal and long-lasting agreements. Conceptually, the ‘team’ sub-structures we create may seem closer to the formal ‘teams’ Munyon references [117]; however they may also represent those intra-organisational groups that form more organically with a social component (i.e. ‘coalitions’), in the same way that hierarchies can emerge in an informal fashion. In either case, later uses of the term *'team'* in this chapter will refer to the local sub-structures we create within the hierarchical k-ary tree networks.

5.3 Extending our model: incorporating status

In Chapter 3, we used homophily as an element of *strength* in our model inspired by Latané's social impact theory. Here, we extend this *strength* to also incorporate *status*. Where previously strength was simply the similarity between two agents, we now multiply this similarity by the edge weight between the nodes in question. This edge weight represents status. Where the similarity between two agents i and j has previously been defined (see Section 2.3.1) as:

$$sim_{i,j} = \frac{1}{F}\sum_{k=1}^{F}\delta_{\sigma_k^i,\sigma_k^j}$$

then $strength_{ij} = sim_{ij} \times w_{ij}$ where w_{ij} is the edge weight from node i to j.

Therefore the trait-scoring formula introduced in Section 3.2 has now been extended to:

$$ts_{\alpha,k,i} = \sum_{j \in N^+(i)} strength_{ij}\delta_{\sigma_k^j,\alpha}$$

where $N^+(i)$ is the set of out-neighbours from node i (i.e., those i copies *from*).

5.4 Methodology

In this section we explain how our model is applied in the context of possible organisational structures, with particular reference to hierarchies and embedding of teams.

5.4.1 Tall vs flat k-ary trees

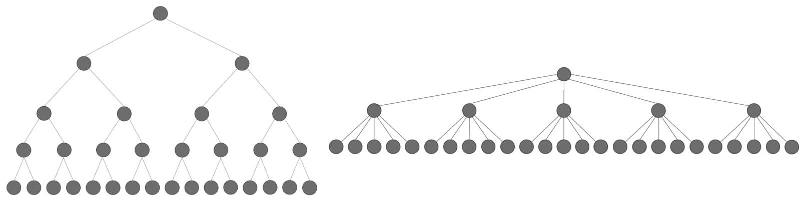

Figure 5.1: A taller binary tree, $k = 2, h = 4$ (left), and a flatter tree $k = 5, h = 2$ (right). Image from [113].

To model the effects of tall-vs-flat hierarchies, we place our agents in balanced k-ary trees of height h and branching factor k. Each tree has a root at level 0 with k children arranged in h levels of descendants. Wider, flatter trees have low h and high k; deeper narrower trees have higher h and lower k (Figure 5.1). Values for h and k were chosen to maintain a relatively similar number of agents across different tree structures (between 1023 and 1555, see Table 5.1). For the purposes of examining the effects of taller-vs-flatter trees, we initially only use undirected edges.

Table 5.1: Tree structure configurations as networks, table from [113].

Branching factor (k)	Height (h)	# of Agents	# of (undirected) Edges
2	9	1023	1022
3	6	1093	1092
4	5	1365	1364
6	4	1555	1554
10	3	1111	1110

We choose number of features, F and traits, q based on values which exemplify the behaviours found in the model, between the states of complete mono-culture and complete polarisation. We run our model on $F = 5, q = 3, 5, 10, 15, 20, 25, 50$ and $q = 10, F = 3, 5, 10, 15, 25, 50$. For each parameter set, we run the simulation 20 times,

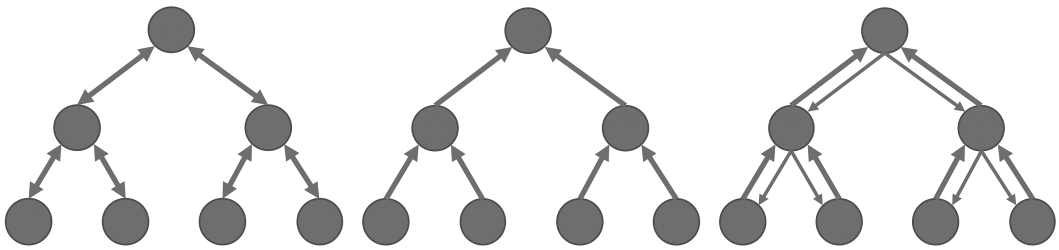

(a) Undirected (b) Exclusively top-down (c) Asymmetric bi-directional

Figure 5.2: Different trait-copying scenarios within tree structures. In undirected trees (5.2a) linked nodes may copy traits from each other in either direction. In a scenario where influence is exclusively top-down (5.2b), edges are directed such that nodes may only copy from their superior - thus traits (and influence) flow downwards. In the asymmetric bi-directional scenario (5.2c), edges exist in both directions but typically have a greater weight where a subordinate copies from their superior. Figure from [114].

each with a different random seed and randomly assigned starting traits. We also run Axelrod's model on similar trees for comparison.

5.4.2 Hierarchical status

Where Section 5.4.1 used undirected edges in tree structures, we now also apply directed and weighted edges to simulate *hierarchical status* (as described in Section 5.2). Using an extended version of our model given in 5.3, we simulate 'top-down' hierarchical status by increasing the weight on directed edges that allow copying from a node's superior. We use our directed and weighted edges in three ways: undirected, exclusively top-down influence, and asymmetric bidirectional influence (Figure 5.2). In the following sub-sections we outline the different scenarios for influence between nodes.

Undirected edges: no hierarchical status

Where edges are undirected and unweighted, hierarchical status is essentially absent. Influence flows both ways along a connection, and subordinates may influence their superiors with as equal a weight as vice-versa (Figure 5.2a). While nodes higher in the tree may have greater centrality and lower average shortest paths lengths to the rest of the network, they are affected by an equal amount of influencing agents to those over which they have influence.

Exclusively top-down influence

Where influence is exclusively top-down, agents may only copy from those directly above them, and only be copied from by those directly below. Edges are directed in such a way that nodes have only out-edges to their superiors and in-edges from their subordinates (Figure 5.2b). Aside from those at the very top and bottom of the tree, each agent may copy from one other and be copied from by k others.

Asymmetric bi-directional influence

We extend the top-down structures by adding additional directed edges allowing upward influence. Agents can now copy both from above and below. However, unlike in the undirected scenario, upward and downward edges may now have different weights. A superior will likely exert more influence upon a subordinate than it receives in return (Figure 5.2c).

5.4.3 Teams

To generate the *team* structures described in Section 5.2.2, we define a team in our network as a set of all nodes containing the same 'parent' node in the tree. In addition to the 'vertical' links in the k-ary tree (i.e. edges to superiors/subordinates), our teams will introduce horizontal links between members of a team.

First, we create k-ary trees (as in Section 5.4.1). Then, each node is assigned to a *team*. The nodes in our tree are numbered sequentially, from node 0 at the top to the last node ($id = |A| - 1$) at the bottom. Each team is numbered by its parent; for example, those agents which are child-nodes of node 3 are assigned to team 3. Node 0 is a special case; with no peers or superiors it sits alone at the top of the tree, a member of no team. Every other node can belong to one team only, and the size of each team is k.

Within each team, for each possible link between members we create an edge with probability p. Where $p = 1$, every team would form a complete subgraph. For example, where $k = 4$ and $p = 1$, each team in the structure would be fully connected, with 6 edges (in the undirected case) linking 4 nodes. Where $k = 4$ and $p = 0.5$, each team would have on average 3 edges randomly distributed across the 4 nodes.

In directed graphs, there are two potential edges between each pair of nodes within a team: inward and outward. An edge created from one agent to another may not necessarily be reciprocated. The edge weight between members of a team is always set to 1, but the weight between team members and superiors/subordinates may be varied to simulate hierarchical status. As we did with basic tree structures in Section 5.4.2, we examine three different scenarios of hierarchical status:

- Undirected team structures: nodes may copy from their team-mates, subordinates and superiors with equal weight (Figure 5.3a).
- Exclusively 'top-down' influence: Nodes may only copy from their peers and superiors (Figure 5.3b).
- Asymmetric bi-directional influence: Nodes may copy from their team-mates, superiors and subordinates; but typically with greater weight from their superiors. (Figure 5.3c).

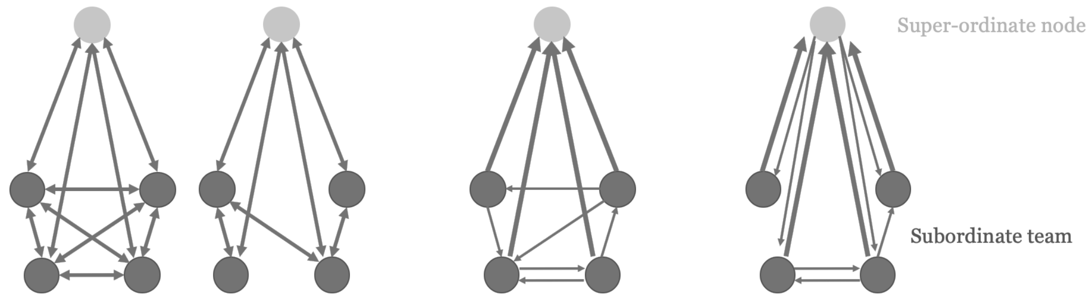

(a) Undirected (b) Exclusively top-down (c) Asymmetric bi-directional

Figure 5.3: Different trait-copying scenarios within team structures. We take the same undirected and directed structures used for trees (Figure 5.2) and add intra-team edges with probability p. Figure 5.3a depicts two undirected teams of size $k = 4$, where the left team is fully connected i.e. $p = 1$. The team on the right has $p = 0.5$, and therefore only 50% of possible edges have been created. Figure 5.3b depicts a scenario where upward influence is absent; nodes may only copy from their team-mates and superior; never from their subordinates. Note that the graph is directed, and that there is a probability of creating an edge in either direction. In this example, $p = 0.5$ and 6 directed edges out of a possible 12 have been generated within the team. Figure 5.3c illustrates the asymmetric bi-directional scenario for $p = 0.25$; nodes may copy from their team-mates, subordinates and superiors but will typically copy with greater weight from their superior. Figure from [114].

5.4.4 Results metrics

Similarly to previous chapters, we examine the composition of the edges within the stabilised simulations; whether they link completely similar agents (a homogenised edge), completely dissimilar agents (a polarised edge), or somewhere in between (a mixed edge). These edges are described in more detail in Section 3.3.2.

As we are studying hierarchies and status, we are also interested in ability of higher nodes to influence others. To do this, we record the topmost (root) node's starting traits at the beginning of the simulation. When the simulation reaches stabilisation, we compare the original traits of the root node with others across the structure. We consider the nodes at the top of the structure to have had a large influence on the organisational culture where their traits are widely adopted by the rest of the population. The average similarity of two nodes with randomly assigned traits is $\frac{1}{q}$. We can use this as a baseline; if the topmost node only has $\frac{1}{q}$ of its traits adopted by others, it has not had any more influence than any other node on average.

We also compare the topmost node's traits at stabilisation to its original starting traits. Should the root node retain a large proportion of its starting traits, then it has not been greatly influenced by the rest of the population.

5.5 Results

We examine the network clustering and path-length that result in the generated k-ary trees and team structures (sec:hierarchiesResultsNetwork). We organise the cultural results by undirected trees (i.e. no hierarchical status, Section 5.5.2), hierarchical status in trees (Section 5.5.3), and results from team structures (Section 5.5.4).

5.5.1 Network characteristics

K-ary trees represent an extreme case where clustering is zero; none of a node's neighbours will be connected to each other. The average node-degree of the basic k-ary trees also remains similar (close to 2) across tall and flat structures as the increase in subordinates (k) in a flatter tree is balanced by the greater number of nodes at the lowest level, who all have degree 1. As the horizontal links of teams are introduced, clustering increases with p. The increase of clustering with p is amplified as tree structures get flatter and the teams within them consequently get larger (Figure 5.4a).

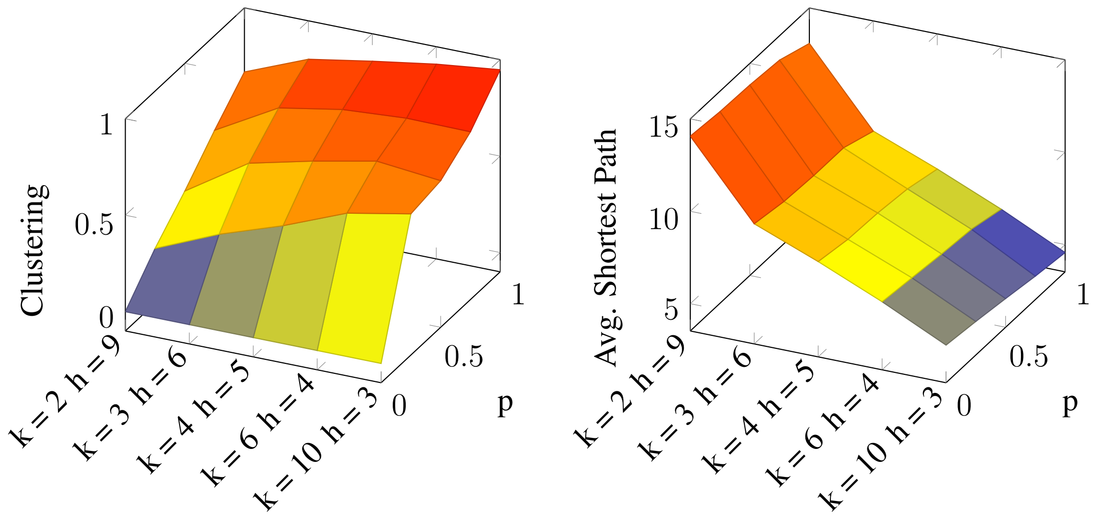

(a) Clustering coefficient in trees and teams

(b) Average shortest path in trees and teams

Figure 5.4: Network characteristics of trees and teams. For all basic k-ary trees, clustering is 0. Clustering is greater in flatter structures that contain teams, and increasing intra-team connectivity p greatly increases clustering. Flatter trees also have much lower average path lengths, as there are less levels of hierarchy to navigate. Increasing p has little effect on path length.

The average shortest path is significantly lower for flatter trees vs tall; the increased hierarchical layers of the latter cause greater distance between the top and bottom of the tree. Introducing teams only has a minor effect in reducing this path length, even where $p = 1$ (Figure 5.4b)).

5.5.2 Undirected taller vs flatter trees

When edges are undirected, flatter trees tend to show greater convergence than taller, narrower trees. This manifests itself as a smaller overall number of cultures (Figure 5.5a) and a greater proportion of homogenised edges (Figure 5.6b). In contrast to flatter trees, taller and narrower structures instead exhibit greater cultural mixing (Figure 5.6c). The proportion of polarised edges appears to be similar across different tree structures (Figure 5.6a). As the number of features F is increased, the difference between tree structures grows (figures 5.5a, 5.6). However, varying traits did not seem to cause much distinction between the results of the different hierarchical structures (Figure 5.5b).

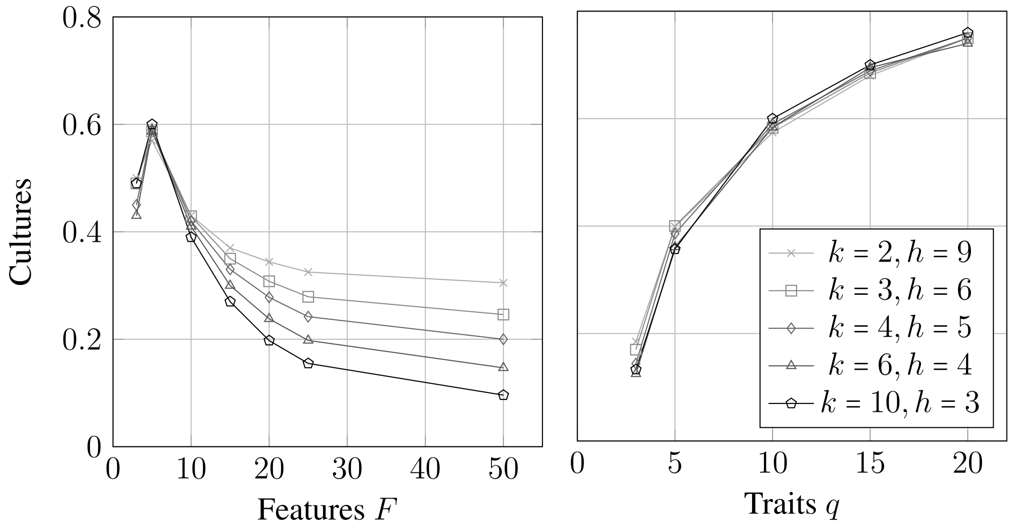

(a) (b)

Figure 5.5: Number of cultures in tall-vs-flat trees. In Figure 5.5a, we vary the number of features, F. The number of cultures decreases as the size of the features vector, F, is increased. Deeper trees maintain a higher number of cultures than flatter trees. Note: $q = 10$. Where $q = 5$ and $q = 15$ similar tall vs flat patterns are present. In Figure 5.5b, we vary the number of traits q. The number of cultures increases as the number of possible traits, q, increases; however there is no discernible difference between different tree structures for F = 5. Figure from [113].

5.5.3 Hierarchical status in taller vs flatter trees

Exclusively top-down influence

In simple k-ary trees, nodes only have connections to their superiors and subordinates. As upward influence is absent, all nodes are completely and exclusively influenced by

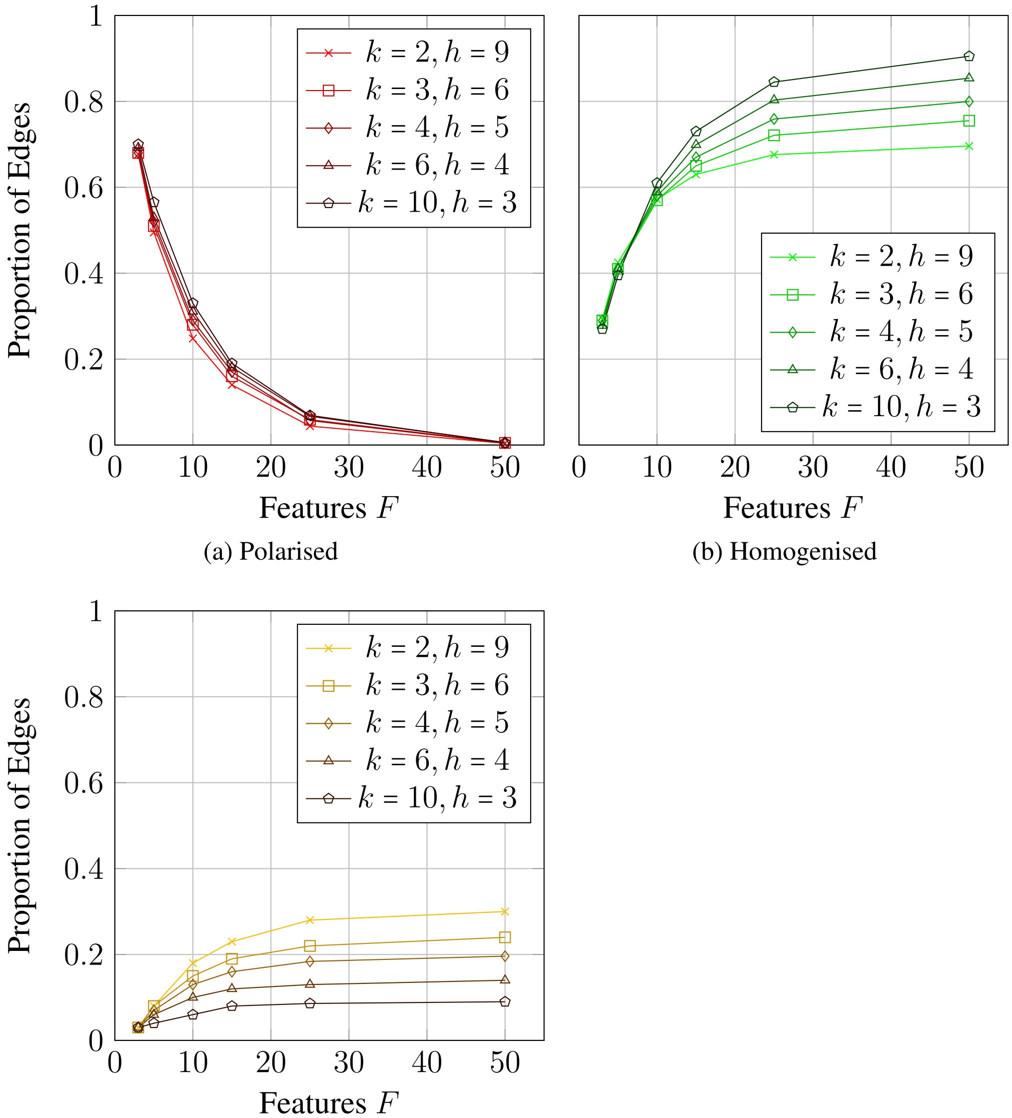

(a) Polarised

(b) Homogenised

(c) Mixed

Figure 5.6: Proportions of edge types for tall-vs-flat trees. As the size of the features vector, F, increases, the number of polarised links drops (5.6a). These links instead become homogenised or mixed. The amount of cultural mixing is greater for deeper trees (5.6c) whereas flatter trees show greater convergence (5.6b).

their superior. This influence from only one other node renders any status multiplier on edge weight moot; so too is compound influence. The ‘cultural mixing’ that is usually a key characteristic of our model cannot occur at stabilisation; the simulation will eventually reach a stable state where all agents are either completely the same as

Table 5.2: A comparison between edge statuses in trees where influence is exclusively top-down, and edge statuses in trees where edges are undirected. Unlike in cases where influence is bi-directional, the simplistic structure of k-ary trees ensures that no 'mixed' links can remain when trait-copying occurs only from parent to child. Table from [114].

	Exclusively top-down influence			**Undirected graph**		
	Polarised links	Homogenised links	Mixed links	Polarised links	Homogenised links	Mixed links
$k = 2, h = 9$	0.056	0.943	0	0.14	0.628	0.232
$k = 3, h = 6$	0.068	0.932	0	0.164	0.648	0.188
$k = 4, h = 5$	0.084	0.916	0	0.17	0.673	0.157
$k = 6, h = 4$	0.095	0.905	0	0.184	0.7	0.116
$k = 10, h = 3$	0.113	0.887	0	0.195	0.73	0.075

their direct superiors or completely dissimilar (see Table 5.2). If a node begins with even a single randomly assigned trait in common with its superior, it will eventually become identical to its superior. Global convergence is extremely high but not total, as there exist instances where agents have no traits in common with their superior and cannot copy from them, essentially insulating that sub-tree from any dominant cultures hegemonising the majority of the structure. In contrast to undirected trees, taller trees show slightly greater homogenisation.

Asymmetric bidirectional influence

The introduction of even a small amount of influence counter to top-down hierarchical status can have a notable effect. In a taller tree ($k = 2, h = 9$) where there is some upward influence, convergence only reaches similar levels to that of the exclusively top-down scenario when the ratio is around 40:1 of top-down vs bottom-up influence (Figure 5.7). At this ratio and above, 'mixed' links cease to be present and the number of polarised and homogenised links are close to that of the exclusively top-down simulation. In flatter trees, the points at which mixed links disappear occur at greater weights: around 100:1 for a tree of $k = 6, h = 4$.

Despite the results of a high downward-to-upward influence ratio largely matching those where influence is exclusively top-down, some differences remain. Where influence is exclusively top-down, the topmost node unsurprisingly retains all of its original starting traits. However, when there is at least some upward influence the topmost node is less able to totally influence its subordinates (Figure 5.8a). Even where there is a very large ratio of downward to upward influence ($> 50 : 1$), the topmost nodes on average only retain 0.6 of their original traits for deeper trees ($k = 3, h = 6$), and 0.85 for flatter trees ($k = 6, h = 4$).

Correspondingly, the amount of the topmost node's starting traits which are adopted by other agents is lower. At around the same point at which mixed links disappear (Figure 5.7), the degree of similarity between the population and the topmost node's starting traits appears to reach a ceiling above which it doesn't climb; even adding extra weight to downward influence beyond this point will not increase the spread of the topmost node's traits (Figure 5.8b).

5.5.4 Team Structures

Undirected Team Structures

It is perhaps an intuitive expectation that increasing the average degree and connectivity between agents who behave homophilically will likely drive convergence toward greater similarity overall. Indeed, simulations on ring-lattices (Section 4.3.2) showed that increasing node-degree in Axelrod's model will cause greater convergence; lowering the average number of distinct cultures and increasing the proportion of homogenised links. For our model, the results were more nuanced (Section 4.3.2); greater degree caused a lower number of cultures but the proportion of homogenised links

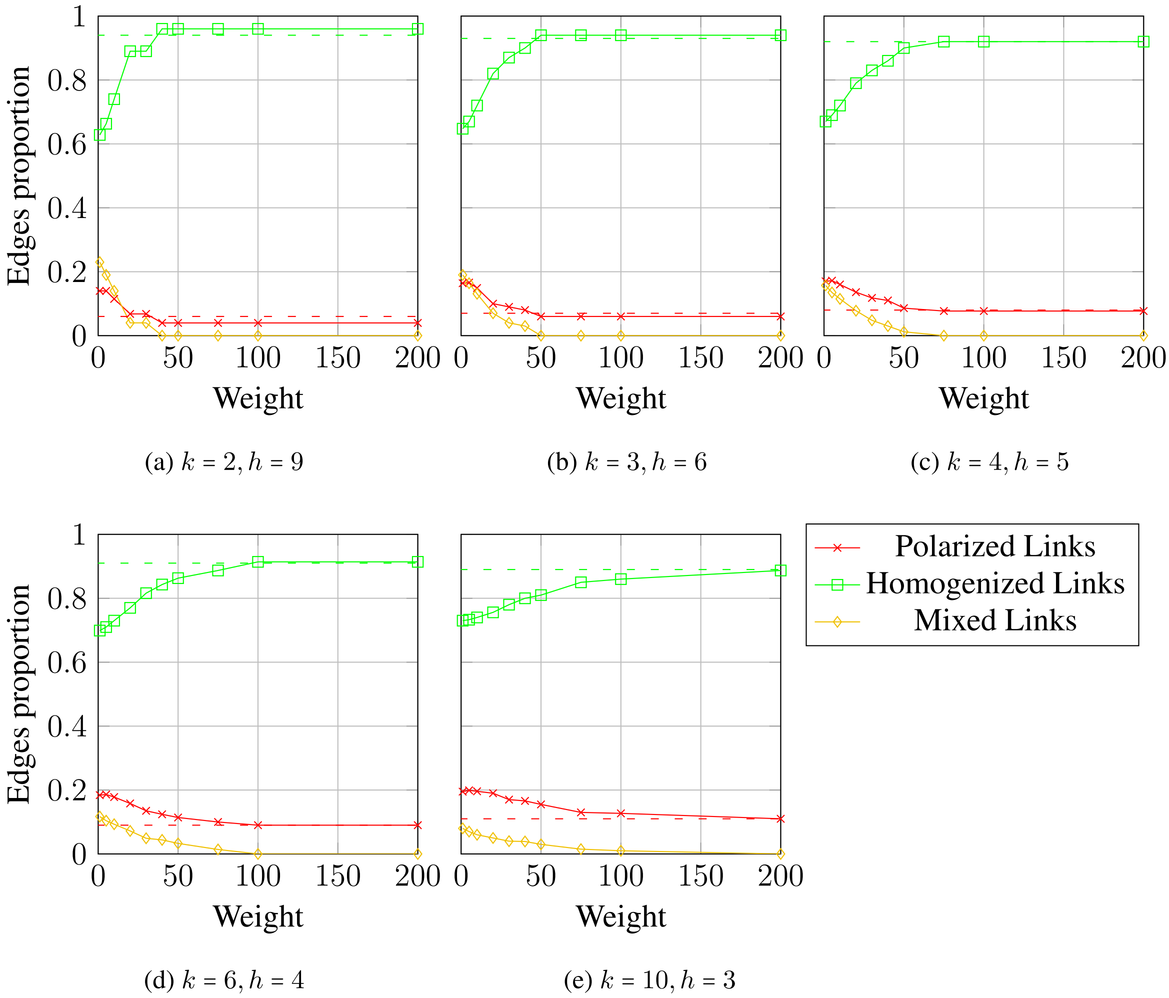

(a) $k = 2, h = 9$ (b) $k = 3, h = 6$ (c) $k = 4, h = 5$

(d) $k = 6, h = 4$ (e) $k = 10, h = 3$

Figure 5.7: The effect of greater ratios of downward vs upward influence from taller (5.7a) through to flatter trees (5.7e). The *weight* (x-axis) represents hierarchical status, and refers to the edge weight w_{ij} applied where subordinates are influenced by their superiors (see Section 5.3). The amount of hierarchical status weight required to raise homogeneity increases as trees get flatter. Mixed links eventually drop to zero. Dashed lines indicate the amount of polarisation/homogenisation when influence is exclusively top-down (Section 5.5.3). As hierarchical status weight rises in relation to upward influence, the proportion of polarisation/homogenisation eventually settles close to that of the exclusively top-down scenario. The point at which this happens is where mixed links reach zero. However, it often requires a large ratio of downward-vs-upward influence to reach this point.

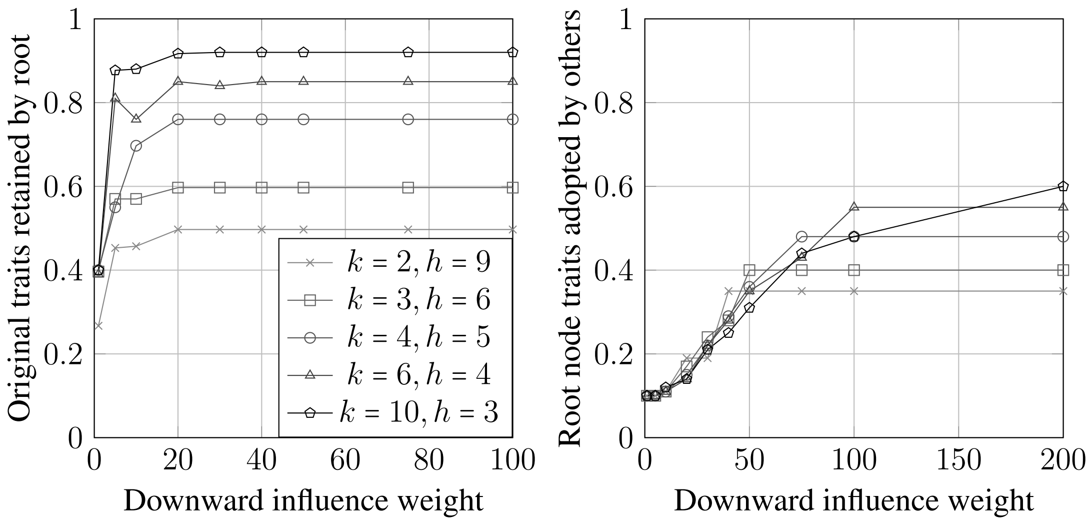

(a) Proportion of topmost node's starting traits which are retained at stabilization

(b) Proportion of topmost node's starting traits that are adopted by other nodes

Figure 5.8: The influence of the topmost node vs upward influence. 5.8a depicts the proportion of starting traits the topmost node retains at the end of the simulation. 5.8b illustrates the proportion of starting traits from the topmost node that are adopted by other agents at the end of the simulation. When there is no upward influence, the topmost node always retains 100% of its starting traits, but this is not the case when there is some upward influence. Even when polarisation and homogenisation approach the levels found in an exclusively downward influence model (Figure 5.7), the topmost node does not exert the same amount of influence as it did where upward influence was absent. The top node does not retain all original traits even when the downward influence is many times that of the upward influence. Additionally, the topmost node's influence appears to be greater in flatter trees.

was largely the same. Results on our 'team' structures show different behaviour still: homogenisation may *drop* when additional links are added.

As the likelihood p of connections between peers within a team increases, the proportion of homogenised links undergoes a u-shaped dip; one which is exaggerated for flatter-shaped organisations (Figure 5.9). Where homogenisation decreases, mixing increases. This is likely another example of the differences caused by influence from multiple sources; as agents are influenced by several others simultaneously, a small increase in their node degree may expose them to a greater number of different traits - but without those traits necessarily becoming 'dominant' locally. Yet further increases in local connections within a group may increase the likelihood of a trait reaching fixation within a team, causing the slight return to convergence as p approaches 1.

Regardless of the value of p, the average similarity of other agents with the topmost node's starting traits remains between 0.098 and 0.102 (where $q = 10$). As the average similarity of any two nodes with randomly assigned traits is 0.1, this indicates that the introduction of intra-team connections has little or no effect on the influence of the topmost node where edges are undirected.

No upward influence

Next we simulated a scenario where no upward influence is possible; nodes may only copy from their peers and superiors. We expected the addition of horizontal links between team members to undermine the influence of the topmost node as p increases: that these extra edges between team-mates would strengthen intra-team homophily and offer 'resistance' to top-down influence. This was not always the case.

When the weight of hierarchical status (downward influence) is low, increasing the number of links p causes less adoption of traits from the top of the organisation, as expected. Where the weight of downward influence is 1, and $p = 1$, the proportion of traits adopted form the topmost node drops to 0.1, the same average proportion of traits in common between two nodes with randomly assigned traits. This means the effective influence of the topmost node is no more than any other node in the structure (Figure 5.10).

However, at higher weights of top-down influence, the introduction of additional links

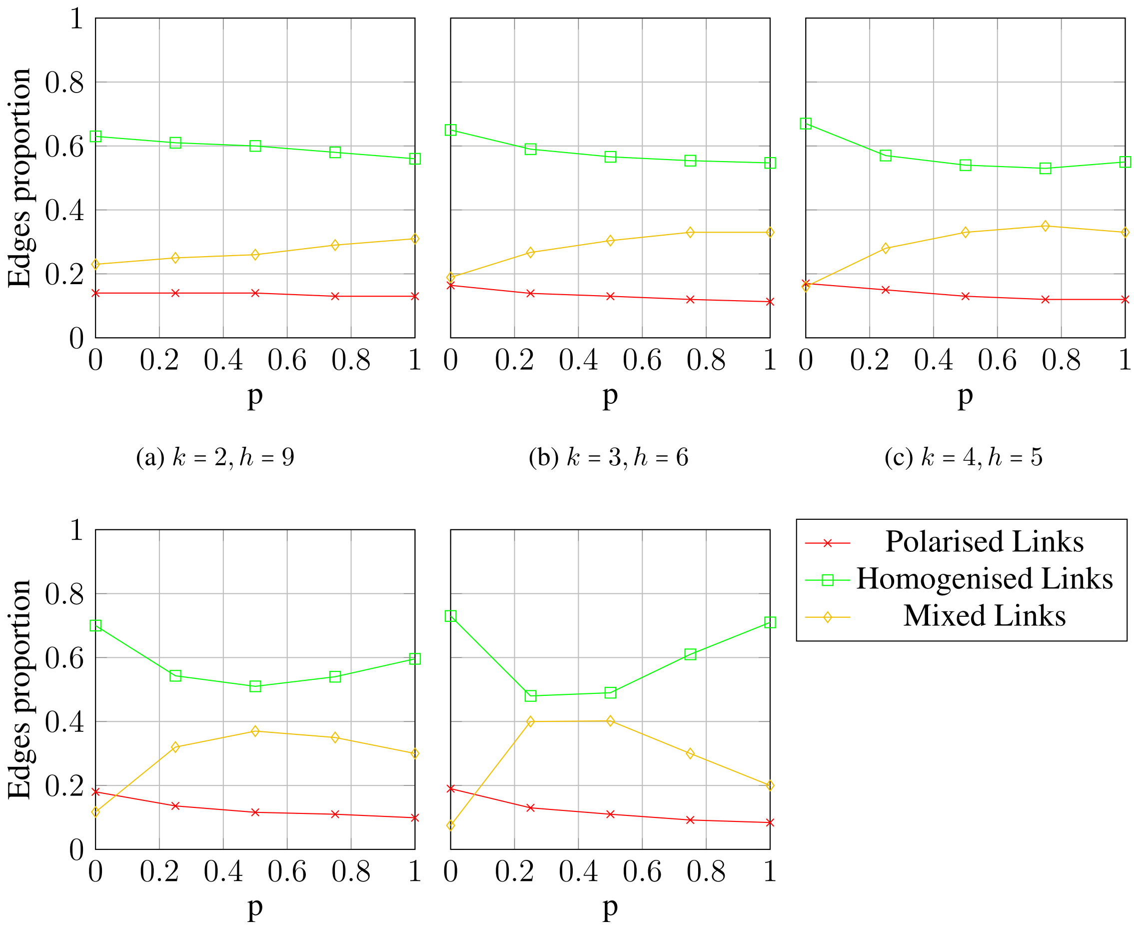

(a) $k = 2, h = 9$ (b) $k = 3, h = 6$ (c) $k = 4, h = 5$

(d) $k = 6, h = 4$ (e) $k = 10, h = 3$

Figure 5.9: The effect of greater intra-team connectivity from undirected taller trees (5.9a) through to flatter trees (5.9e), where p is the probability of an edge existing between two team-mates. Increasing the number of edges between nodes of the same team (p) does not necessarily drive convergence; in a compound influence model increasing connectivity may instead encourage greater cultural pluralism ('mixed links'). Increasing horizontal connectivity in tree structures produces a 'u-shaped' effect on homogenisation.

between peers instead serves to increase the amount of traits adopted from the topmost node. When the weight of influence from higher status nodes is greater, links between their subordinates have a further reinforcing effect in promoting their superior's traits. In addition, the greater connections within larger teams mean that this effect is more

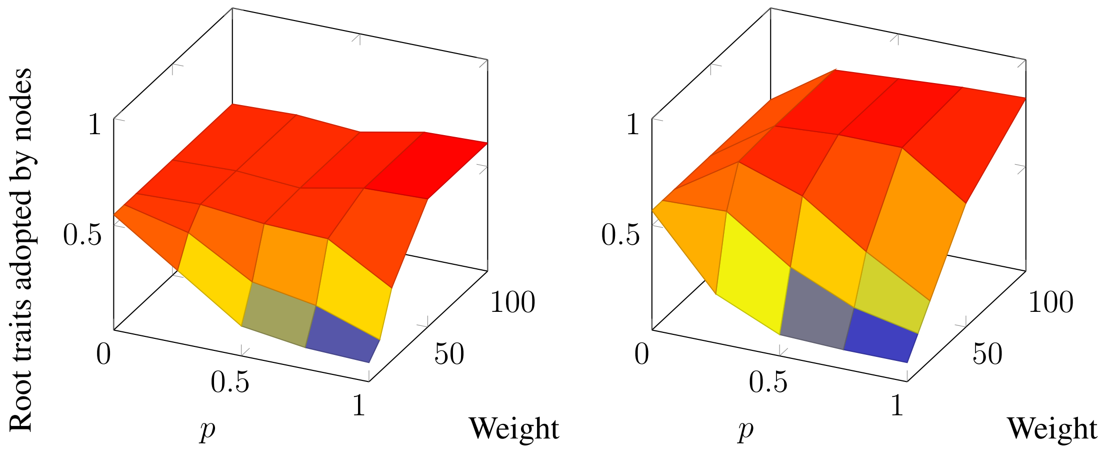

(a) $k = 3, h = 6$

(b) $k = 6, h = 4$

Figure 5.10: Proportion of topmost node's traits adopted by others in team structures. When the 'top-down' hierarchical status weight (y-axis) in the structure is low, increasing the connectivity of teams (p, x-axis) weakens the effective influence of the high level nodes. This is the area of the charts shaded blue/green, a 'valley' where subordinates are influenced little by their superiors. Where $p = 1$ the topmost node traits adopted drop to 0.1, no greater than the average similarity of two nodes with randomly selected traits; at this point the high connectivity of subordinates with each other has reduced the topmost node's influence to the same as any other node. However, when the 'top-down' status weight is increased, the increased connectivity of subordinate teams serves to reinforce the topmost node's effective influence. The red-shaded areas of this chart, where the influence of superiors is highest, occur where there is at least some intra-team (p) connectivity, and often where p is higher. This effect is more pronounced for flatter trees (5.10b) than taller (5.10a).

pronounced in flatter, broader structures (Figure 5.10b) than deeper, narrower organisations (Figure 5.10a).

Asymmetric bi-directional influence

Our expectation was that any effects of upward influence apparent in tree structures would be greater in corresponding ‘teams’ structures. However, where upward influence is present, there are only minor differences between the results from basic k-ary trees (Section 5.5.3) and team structures where teammates are fully connected.

When we examine the number of the topmost node’s traits copied by other agents, we find little difference between the results from basic bidirectional trees to those that contain fully connected teams. Where small differences exist, there does not seem to be a discernible pattern that would suggest that team structures either increase or decrease root influence across different bidirectional tall or flat structures.

At low levels of downward influence, more connections between teammates causes an increase in the proportion of ‘mixed’ edges - those between agents with non-identical but not dissimilar cultures. As ours is a model of compound influence this is intuitive; the greater a node’s degree the greater the likelihood that it can be influenced simultaneously by multiple competing cultures. However, as the weight of downward influence increases, these differences between basic trees and team structures become smaller or cease to exist.

5.6 Discussion

Where basic undirected k-ary trees with no team sub-structures are used (Section 5.5.2), flatter trees show a clear trend toward homogeneity. This may seem counter-intuitive if one envisions a typically (and perhaps, traditionally) tall hierarchy as more rigid and autocratic. However, the greater degree of intermediate ‘hub’ nodes, social reinforcement from more subordinates, and possibly shorter paths of communication (Figure 5.4b) may aid convergence. Note however, that in these most basic of k-ary trees, increasing k does not increase clustering above zero (Section 5.5.1 and Figure

5.4). It is also interesting to note that despite the shorter path lengths of flatter trees offering a possible explanation for greater convergence, in ring lattices we did not find path length to have a great effect on the results of our model (Figure 4.10, Sections 4.3.2, 4.4).

Where we first introduce hierarchical status to basic k-ary trees as an exclusively top-down phenomenon (Section 5.5.3), the results are in some ways predictable. Where an agent can only be influenced by one other - its superior - there is no way cultural mixing can endure. The simulation ends in a similar state to that of the typical Axelrod model; all homogenised and polarised edges. While a scenario where agents can only be influenced by their direct superior is not impossible, it is likely very rare except in the most authoritarian of environments.

The impact of upward influence is evident from Section 5.5.3. Where superiors allow themselves to be influenced even a small amount by their subordinates, popular traits can spread upwards and a greater mix of cultures can exist in an organisation. There are many instances where such upward spread is considered beneficial. While our model has no notion of 'good' or 'bad' traits, it would generally be considered desirable for any superiors who valued efficiency to allow a successful trait developed at a low level to spread up and across the organisation. In reality, office politics and power struggles may undermine such a notion [29, 153]; our model does not currently model conscious and deliberate 'blocking' strategies, although 'blockages' in the spread of traits do occur in circumstances where superiors and subordinates are completely dissimilar.

The introduction of team structures (Section 5.5.4) further illustrates the nuances in behaviour when trait spread is based on compound influence, rather than the dyadic-copying of models such as Axelrod's or many infection-based models. As we have seen in previous chapters, where agents can be influenced by multiple neighbours simultaneously, cultural mixing can occur. In our k-ary trees with team sub-structures, an increase in intra-team links creates an inverted u-shape in the proportion of mixed links (Figure 5.9). It should be noted that while non-monotonic, at $p = 1$ (a completely connected team) mixing is generally greater than where $p = 0$ (an unconnected team). This is of interest because where in Chapter 4 a decrease in clustering caused an increase in mixing (Figure 4.10), here it is the addition of intra-team links that seems to facilitate cultural mixing, despite the fact that these links increase clustering. Clearly, we cannot predict the amount of homogenisation and cultural mixing from network clustering alone.

A perceived benefit of 'flat' organisations is the supposed reduction of formal hierarchies and their associated politics to better allow power and influence to be more evenly spread across the structure [154, 161]. However, the results illustrated in figures 5.8 and 5.10 suggest that while cultural plurality may increase in flatter structures, so may the superior's influence. Where strongly influential personalities exist in the upper levels of an organisation, flatter structures may further concentrate influence within a small number of these nodes.

As generalisations of hierarchies, basic k-ary trees have obvious limitations. A scenario where an individuals only has contact with their direct superior and subordinates, and these contacts of the ego then have absolutely zero contact with each other, seems highly unrealistic. While the team structures proposed offer a more realistic extension, in some cases the results differ little with the basic trees. This suggests there may be phenomena common to organisational structures such as these. Despite limitations and simplifications in the construction of generalised networks, we believe some of the behaviours observed may be the result of the basic characteristics of a hierarchical structure: typically hierarchies - and most organisations - form a pyramidal shape. High level nodes can cause bottlenecks and blockages to the spread of traits and behaviours up and across an organisation. On the other hand, even highly influential high-level nodes may be swayed by the sheer number of their subordinates adopting a behaviour. In a sense, there is a balance between hierarchical authority and collective influence; the *hierarchical status* introduced in this chapter versus the *number* of simultaneous influences that forms a key part of our model.

We believe the structures used in this chapter may serve as a useful basis for examining the behaviour of other models. Inter-team links may further enhance these structures, but risk diluting the hierarchical nature of the network. A potential future enhancement could be in the form of a multi-layer network, with the formal hierarchy forming one layer and an informal social network forming another - perhaps linking members of different teams or levels. Another approach may be to increase the network distance at which agents may influence each other; we add this notion of *distance* to our model in Chapter 6.

5.7 Conclusion

In this chapter, we have made two main contributions: examining the behaviour of a compound influence model in hierarchical structures and introducing a method of generating and examining 'team' sub-structures within basic k-ary trees. The team structures improve upon basic k-ary trees as a generalisation of hierarchical organisational structures by allowing horizontal interactions where they are most likely: between members of a common team. We believe these structures may be a useful generalisation for other network-based models.

We observe that whether a tree is narrow and tall, or broad and flat, has a tangible effect on the amount of cultural convergence. Also, in support of research question RQ5, we find that where top-down hierarchical status is present, the increased influence of higher level nodes is often counteracted by the number and connectivity of subordinates. A small amount of upward influence may greatly diminish the spread of traits from the top of the hierarchy, even where the weight of top-down status is greater. Increased connectivity within teams of subordinates may diminish the influence of a superior who lacks sufficient status; conversely, a strongly connected team may serve to enhance the influence of a superior where their hierarchical status is high.

The results add further evidence to the notion that the shape of an organisational structure can have a tangible effect on its culture and internal dynamics, in addition to (or even despite) motives and behaviours of individuals. These findings may have implications for how organisational cultures may take hold, and the design of organisational networks and hierarchies in light of their cultural implications. Whether seeking subsidiarity and the development of individual behaviours from below, or a more robust top-down culture, the results from this chapter demonstrate that the design of a hierarchical structure can have a notable effect.

Chapter 6

The influence of social distance

In this chapter, we extend our model of social influence to include *distance*. This allows influence beyond direct connections, to acquaintances and ‘friends-of-friends’. We aim to model effects such as the reach of social influence, and the *immediacy* element of social impact theory [88], using network path-length as an analogy. We introduce two possible mechanisms to capture distance; one based on an inverse-square law, the other on the intermediary paths between agents. Using simulations over network graphs, we explore the effects these processes have when varying clustering and path length in Watts-Strogatz networks. This chapter supports research question RQ6.

6.1 Social Distance

To avoid confusion with everyday terminology, we should differentiate between the *distance* within social networks (either real networks or simulated), and broader metaphorical uses of the term *social distance*. The notion of ‘social distance’ is often associated with Georg Simmel’s essay *‘The Stranger’* [142], who appears to be both near and far to his social group. Owing to academic interest in immigration and racial tensions in the USA in the 1920s, the concept of ‘social distance’ became largely concerned with the differences between individuals or groups [92, 158]. This ‘similarity’ notion of social distance is clearly more closely related to homophily [90, 103], which is already included in our model (see Section 3.1.1). However, Levine and Carter [92] argue Simmel also used ‘nearness’ to describe interactions; the nearness of the stranger being their proximity to, and regular contact with, a social group. Thus a ‘stranger’ in a community may have regular contact with its members, but remain different. This may bear similarities to our scenario in Figure 4.6, where an agent has close ties to others with which it shares little in common. It is the ‘nearness’ of interactions that we aim to add to our model in this chapter.

6.2 The role of Immediacy in Social Impact Theory

In Chapter 3 we focused on the roles of *Strength* and *Number* in Social Impact Theory [88] as inspiration for our extension of Axelrod’s culture model. By allowing multiple sources of influence to act upon an agent, we sought to model influence more social than interpersonal [59], while making a thematic link to the *Number* in Latané’s theory of Social Impact. For *Strength*, we used the existing homophily present in Axelrod’s model, while also using edge weights for potential additional factors of ‘strength’ such as hierarchical status in Chapter 5. Latané describes *Immediacy* as *“closeness in space or time and absence of intervening barriers or filters”* [88]. Proximity in space or time seems similar in spirit to Simmel’s notion of ‘nearness’ (Section 6.1), that those who are ‘near’ would be close geographically, or at least be in regular contact. The “absence of intervening barriers or filters” suggests direct contact between agents, with little or no noise (in the signal sense) or interference from others. In [122], Nowak et al model social impact theory as a cellular automaton, and immediacy as the Euclidean distance between cells. Although our model is based on network graph structures rather than cellular sites, we use a similar principle in that the immediacy of influence between our agents is based on the distance of path lengths between them in the network. Those agents that are a shorter network distance from each other will represent agents nearer in “space or time”; the fewer network hops between them the “absence of intervening barriers or filters”. Accordingly, influence from others further away will likely diminish as distance increases, and intermediaries may act as barriers or filters. The latter forms a key element of the path-based distance method we introduce into our model in this chapter; that nodes along a path may potentially enhance, diminish, or block entirely influence from more distant agents.

6.3 Extending our model: incorporating distance

To add the notion of 'distance' (or 'immediacy') to our model, we base our algorithm on the established assumption that a force such as influence is likely to decay as distance increases [88, 89, 103, 122, 124, 166]. We study two approaches; one based on an inverse-square law (Section 6.3.1), and another more novel process of path-based influence over distance (Section 6.3.2). In either of these methods, an agent being influenced by others at greater distance would seem to be influenced *indirectly* by those more distant nodes. Depending on context this may well be a fitting representation; agents may receive second-hand information and be influenced by the opinions and norms of the wider population.

However, this second-order influence does not necessarily need to represent a lack of contact between nodes at distance, but rather *less immediate* contact. If those vertices in a node's immediate neighbourhood represent its closest contacts, perhaps family members or close friends, those at greater distances may represent colleagues through to acquaintances, those we know by name, those we recognise and finally those with whom we have no contact at all. There are parallels with different group sizes and layers in anthropology, such as families relative to tribes [28]; or support cliques and sympathy groups in Dunbar's work [75]. Agents will generally have their greatest frequency of contact with their closest social circle.

There is also a body of work [27, 40, 58, 69, 133] that proposes that should an individual A have strong ties to two others, B and C, then those others will usually have at least some form of link to each other. Mark Granovetter went so far as to call a structure with an absence of tie between B and C a 'forbidden triad'; asserting that a 'weak

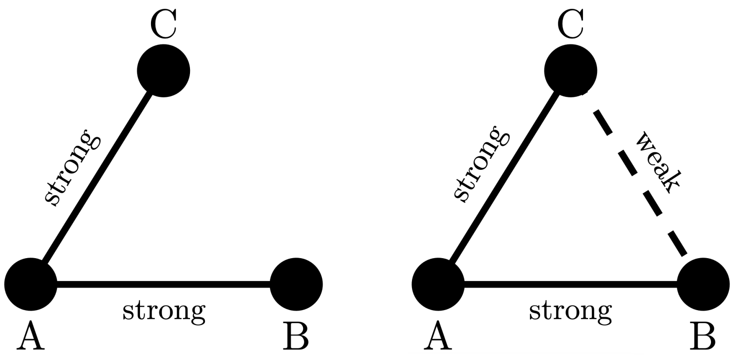

(a) A 'forbidden triad'

(b) A 'forbidden triad' closed with a weak link

Figure 6.1: Granovetter's 'forbidden triad': where strong ties exist from A to B and A to C, some form of relationship must exist between B and C - even if only a weak or indirect one.

tie' will almost always exist to close the triangle [69] (Figure 6.1). The inclusion of weaker influence of nodes at $l > 1$ in effect models such a triadic closure; if strong ties exist as first-order contacts between $A \leftrightarrow B$ and $A \leftrightarrow C$, then B may be influenced by C at distance: a weaker tie. B is likely to have at least some contact with C, albeit with less immediacy.

'Distance' in Axelrod's *The Dissemination of Culture*

It should be noted that in Robert Axelrod's original paper, *The Dissemination of Culture* [6], he refers to interactions over *distance*. For these interactions, agents are able to copy from those beyond their local Von Neumann neighbourhood (i.e. at greater Manhattan distances [22]). While this may make some sense in a cellular automaton, in a network graph the ability of two nodes to interact directly, and with no barriers, interference or noise, means they are connected. If an edge exists between A and B, and A can interact with C under the same conditions, conceptually there is an edge $A \leftrightarrow C$. Therefore the notion of distance described in [6] is really one of connectivity; when 'distance' in the grid is increased the average node-degree increases. This is ac-

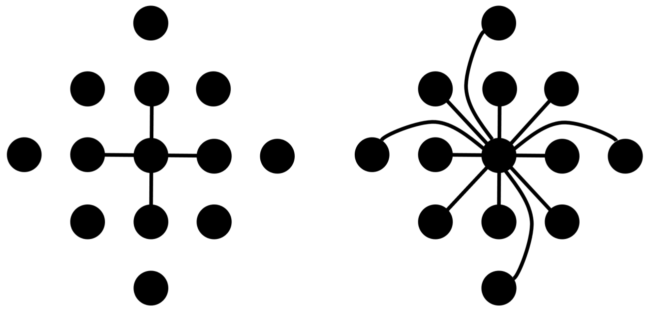

Figure 6.2: 'Distance' in *The Dissemination of culture* can instead be thought of as connectivity. Where the Manhattan distance increases (right), agents create first-order connections to other nodes with no decaying effect of distance.

knowledged in the original paper by the use of the terms 'Neighbourhood' and 'Range of Interaction'. We do not use this characterisation of 'distance', as we have already examined the effects of greater degree and connectivity in Chapter 4.

6.3.1 Diminishing influence over distance as an inverse-square law

As many social simulations and models take inspiration from physics and statistical mechanics [14, 34], several base their modelling of distance decay on concepts from the physical sciences. Newton's law of gravity serves as an analogy for several studies that model some force diminishing over distance [62, 77, 134]. Accordingly, these forces are often modelled as inversely proportional to the square of the distance [45, 122]. In addition, some social phenomena over distance have been found to follow patterns similar to inverse-square functions [89, 149, 166].

We apply an inverse-square function to distance in our model. Distance is modelled as the number of 'hops' between vertices. For example, where an edge $\{a, b\}$ between two nodes exists, the distance between a and b would be 1. Should edges $\{a, b\}$ and $\{b, c\}$ exist, but no edge $\{a, c\}$, then the distance between a and c would be 2. Where l is the distance between node i and j, j's influence over i will be $\frac{1}{l^2} \times influence$.

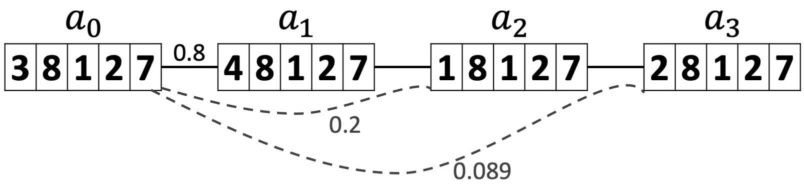

Figure 6.3: Distance as an inverse-square law: As nodes increase in distance from a_0, their influence decays. a_0 has similarity of 0.8 with each of the other agents illustrated. However, as the inverse-square law is applied, a_2 has its trait scores adjusted by $\frac{1}{2^2}$, giving $0.8 \times 0.25 = 0.2$. a_3 has its influence reduced even further: multiplied by $\frac{1}{3^2}$.

The full equation for calculating a trait's score[1] when an inverse-square distance law is applied is:

$$ts_{\alpha,k,i} = \sum_{l=1}^{d} \sum_{j \in A^l(i)} \frac{1}{l^2} sim_{i,j} \delta_{\sigma_k^j,\alpha}$$

where d is maximum distance; j is a node in the set $A^l(i)$ which is the set of all nodes where the shortest path is l hops from i; $sim_{i,j}$ is the similarity between i and j (see Section 2.3.1); $\delta_{\sigma_k^j,\alpha}$ is Kronecker's delta (1 if trait α matches j's k^{th} feature, 0 otherwise). Therefore, where j is within a distance d of i, and holds the trait currently being scored, that trait's score will increase. When all traits from all nodes within distance have been examined, the highest scoring trait will be adopted.

Note that pre-defined edge weights (such as those used in Chapter 5) are not taken into account, as agents may interact with others not in their immediate neighbourhood.

6.3.2 Distance as indirect influence

We have also developed a method of influence over distance, first introduced in [112], where we use the product of influence along a path. Where maximum influence dis-

[1] See Section 3.2 for an explanation of trait scoring.

tance is two, and there exists links $A \leftrightarrow B$ and $A \leftrightarrow C$, C's influence on B is the influence of $A \leftrightarrow B$ multipled by the influence of $A \leftrightarrow C$. In [112] this was expressed more simply as the product of the edge weights along a path, where edge weights were simply the similarity of the edge's two vertices. However, we have since treated edge weight as a separate parameter to allow it to represent other factors such as status (see Chapter 5). Treating the similarity of two agents differently to the edge weight between them also allows us to more easily potentially extend our model to more dynamic network structures without having to recalculate edge weights after every interaction or network change. The updated trait-scoring equation using this method is:

$$ts_{\alpha,k,i} = \sum_{l=1}^{d} \sum_{P \in Paths_i^l} \left(\prod_{n=0}^{|P|-1} strength_{P(n),P(n+1)} \right) \delta_{\sigma_k^{P(l)},\alpha}$$

where d is maximum distance; P is a set of vertices forming a path in the set $Paths_i^l$, which is the set of all possible paths (with no cycles) from i of length l; $strength$ is $similarity \times weight$ (see Section 5.3). Thus $P(l)$ represents the final node on a path, and must hold trait α for α's score to increase. If we expand the full calculations for $strength$ and $similarity$, the equivalent and final formulation of our model of trait-scoring used in this thesis is:

$$ts_{\alpha,k,i} = \sum_{l=1}^{d} \sum_{P \in Paths_i^l} \left(\prod_{n=0}^{|P|-1} \left(\frac{1}{F} \sum_{x=1}^{F} \delta_{\sigma_x(P(n)),\sigma_x(P(n+1))} \right) \times w_{P(n),P(n+1)} \right) \delta_{\sigma_k^{P(l)},\alpha}$$

where w is weight. This expression of the formula is included for completeness, the simplified version above is generally preferred.

In this way, the similarity of intermediary agents along a path is important. Where similarity is high, influence over distance will decay at a lower rate. Where similarity is low, distance decay will be much greater; a polarised edge along the path will block influence entirely. An example of this is given in Figure 6.4. An agent's closest neighbours mediate the influence of more distant contacts, in the way one's closest

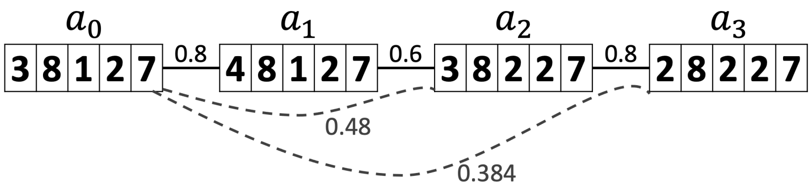

Figure 6.4: Distance as indirect influence: Unlike the inverse-square approach, in our path-based distance method the similarities along the path to a node are also important. While a_0 has similarity of 0.8 with a_2, its influence from a_2 is instead calculated by multiplying the similarities along the path a_0, a_1, a_2, giving $0.8 \times 0.6 = 0.48$.

family may shape one's views of others in the wider community or society. It should also be noted that multiple paths may exist to another node, particularly those closest, potentially raising their influence further.

6.4 Simulations of our multiple-influence model with the additional effect of influence over distance

6.4.1 Methodology

To examine the effects of influence over distance, we have largely repeated the experiments of Chapter 4, but with the extension of our model with the distance methods described above. With both our *path-based* method, and an *inverse-square* law, we examine the effects of increasing influence over distance while also varying homogeneous degree within a ring lattice. We then apply both methods to lattices increasingly rewired into 'small-world' networks, to observe the effects of decreasing global average path length and network clustering.

A list of parameters used is given in Table 6.1. Each parameter set was run 10 times with different random seeds. While the number of runs is small, this is consistent with the methods used in Chapter 4, and is again due to limited resources. The parameters in Table 6.1 give a total of 1260 runs, many of which take several hours to complete. The exponential growth in complexity with the increase of an agent's social ties or distance of influence is analogous with the 'social brain hypothesis' that social group size is limited by cognitive ability [55].

Model	Our multiple-influence model
Distance Methods	Inverse-square, Path-based
Distance	1, 2, 3
Degree	4, 8, 12
Rewiring probability	0.01, 0.05. 0.1, 0.2, 0.5, 0.75, 1
q	10
F	5

Table 6.1: Parameters used in testing effects of distance.

Hypotheses

We hypothesise that at lower degree, results from both distance methods will show little difference to results where influence is from local neighbourhood only. The additional influencing agents will be small in number, and each only have a very small influence because of the effects of distance decay.

When run over ring lattices of higher node-degree, we expect any effects caused by increased distance to be greater, as the number of nodes at distance l will increase with degree. It is possible that higher node-degree may cause either greater convergence *or* greater cultural mixing. As we have seen in Chapter 4, being exposed to nodes beyond one's local cluster can cause a greater mix of traits to be adopted. On the other hand, where agents are 'caught between' two or more cultures, extra influences introduced by our distance methods may act as 'tie-breakers', pushing an agent into one culture or another and driving convergence.

We would expect any effects of 'small-world' rewiring, as observed in Chapter 4, to be exaggerated by the inclusion of influence over distance. Where a single edge exists as a bridge across areas of the ring lattice, neighbours of a bridging node may now exert influence over distance, albeit weakly.

The two distance methods, path-based influence and inverse-square, may show slightly different results. In our path-based influence method, dissimilar nodes along a path can effectively 'block' influence (as similarity will be zero); the inverse-square method will still allow influence from nodes-at-distance who are insulated from the influenced agent by dissimilar nodes.

6.4.2 Results

A: The effect of inverse-square-based distance and node-degree in simple ring-lattices

When node-degree is low ($degree = 4$), adding the process of inverse-square-based distance has only a minor effect. Convergence is slightly higher, as evidenced by the decrease in cultures (Figure 6.5a) and increase in homogenised edges (Figure 6.5c). The number of 'mixed' edges shows almost no change (Figure 6.5d).

We had expected an increase in degree to increase any effects caused by increased distance. Instead, as degree increases the additional influence of distance appears to diminish; results become slightly closer to that of $d = 1$ (Figure 6.5).

B: The effect of path-based distance and node-degree in simple ring-lattices

The effects of applying our path-based distance process, even at low degree, are more pronounced than with the inverse-square-based method, particularly for the number of 'mixed' edges. When $degree = 4$ and $d > 1$ there is greater evidence of convergence shown by an increase in homogenised edges (Figure 6.6c). The amount of cultural mixing is significantly lower when using path-based distance (Figure 6.6d), in contrast to our hypothesis that distance would have little effect at low degree. As degree increases, results for $d > 1$ show a greater increase in homogenised links and convergence.

C: The effect of path-based distance in Watts-Strogatz small-world networks

When we introduce rewired 'small-world' networks, at low degree ($degree$ = 4) the results patterns are largely consistent with those where no distance-mechanic is applied. The differences caused by distance (e.g. greater homogenisation) that we outline above for ring lattices, remain at similar proportions as the lattices are rewired (Figure A.1, appendix).

At higher degrees, the results where path-based distance is applied start to diverge from those where it is not (Figure 6.7). Where the probability of rewiring p is greater than 0.1 a reversal in the decline of homogeneity occurs. The number of homogenised edges rises sharply; mixed links and overall number of distinct cultures appear to drop as the topology approaches that of a random network. These effects are exaggerated by increased degree; at $degree$ = 12 the rise in homogeneity occurs at both lower p and lower distance than results from $degree$ = 8 (Figure A.2, appendix). However, where rewiring probability reaches 1, essentially a random graph, cultural diversity may recover slightly.

To further explore the decline in cultural mixing beyond $p > 0.1$, we examine the make-up of rewired edges versus those that remain part of the ring lattice. Where distance of influence $d = 1$, the portion of rewired edges which link homogeneous cultures grows slightly. However, this growth is not enough to counteract the substantial decline in homogenisation that occurs in the non-rewired edges of the ring lattice (Figure 6.10a). Thus, for $d = 1$, homogenisation is lowered in networks with more rewired edges.

Where distance $d > 1$, the opposite often occurs beyond $p > 0.1$. The portion of

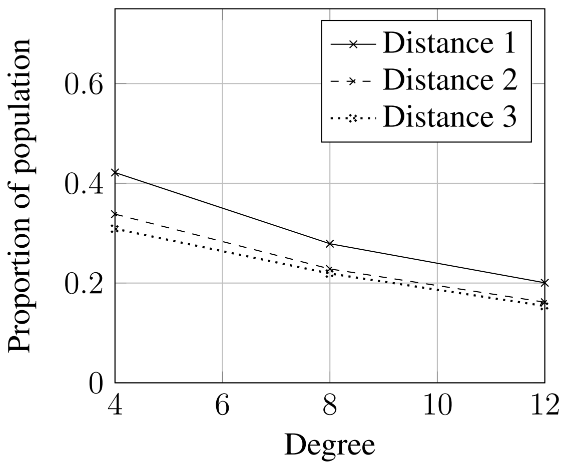

(a) Cultures

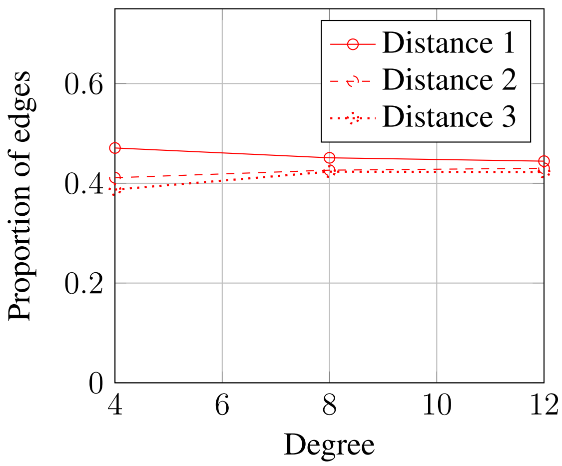

(b) Polarisation

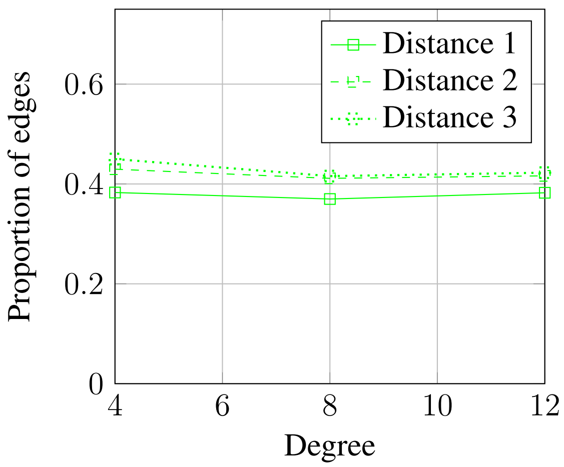

(c) Homogenisation

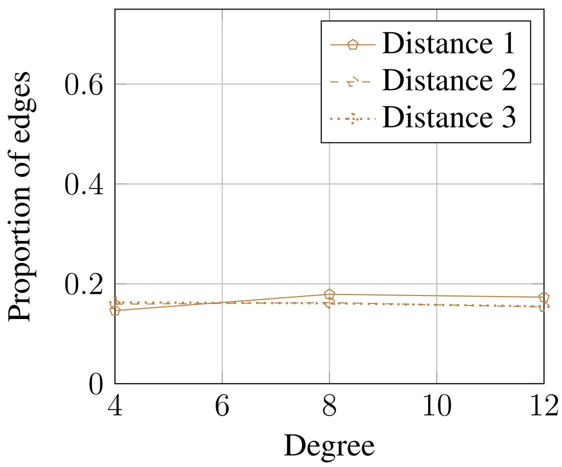

(d) Mixing

Figure 6.5: Results of inverse-square distance over varying degree ring lattices: Applying an inverse-square distance effect causes minor differences to results in ring lattices. Greater cultural convergence is suggested by a lower number of cultures (6.5a), lower polarisation (6.5b), and greater homogenisation (6.5c). As degree increases, there appears to be less difference in results for influence over different distances.

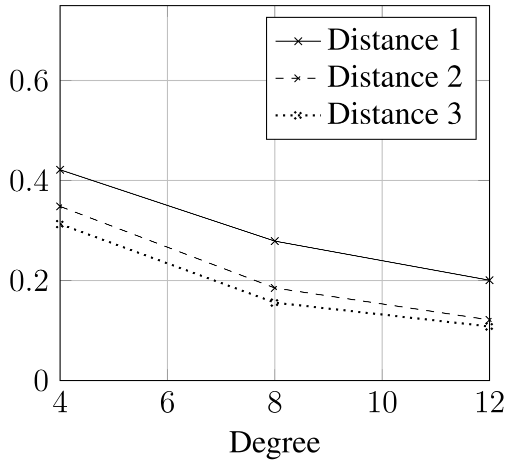

(a) Cultures

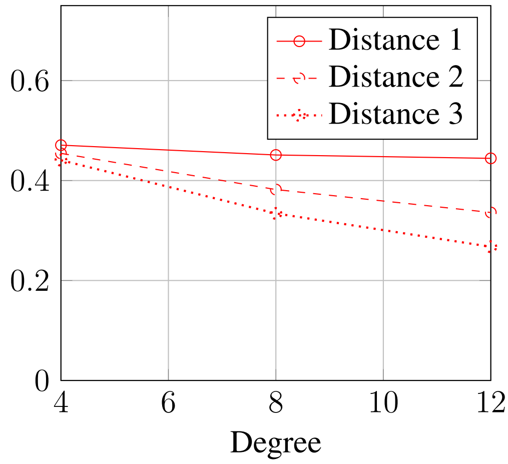

(b) Polarisation

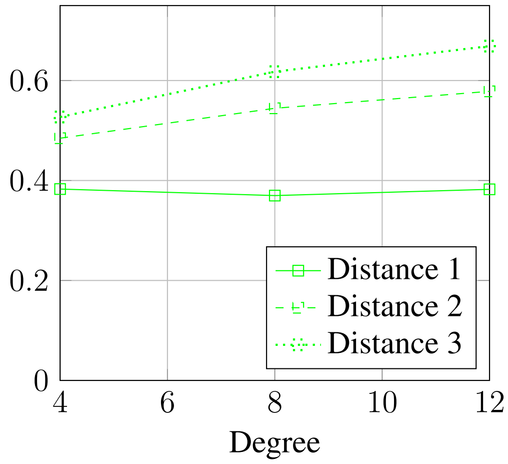

(c) Homogenisation

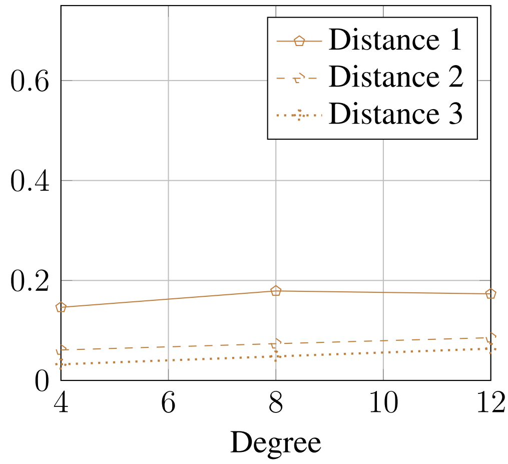

(d) Mixing

Figure 6.6: Results of path-based distance over varying degree ring lattices: When our path-based distance method is applied, the effects are greater than when using an inverse-square method (Figure 6.5). Convergence increases with both increased distance and increased degree (6.6b, 6.6c). Mixing (6.6d) is lower than when using the inverse-square method (Figure 6.5c).

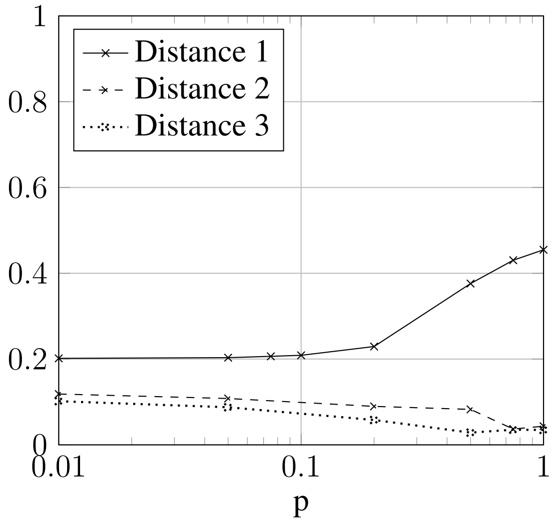

(a) Cultures

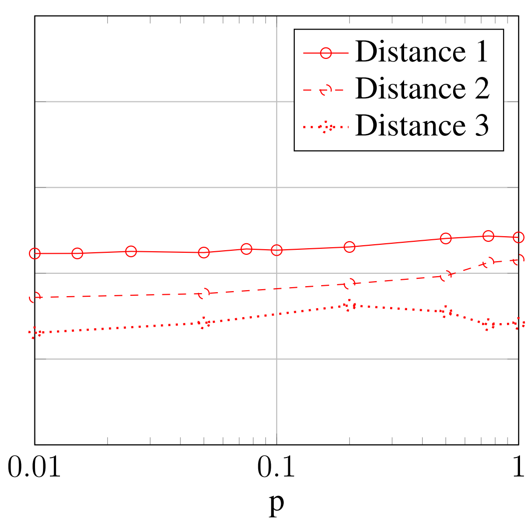

(b) Polarisation

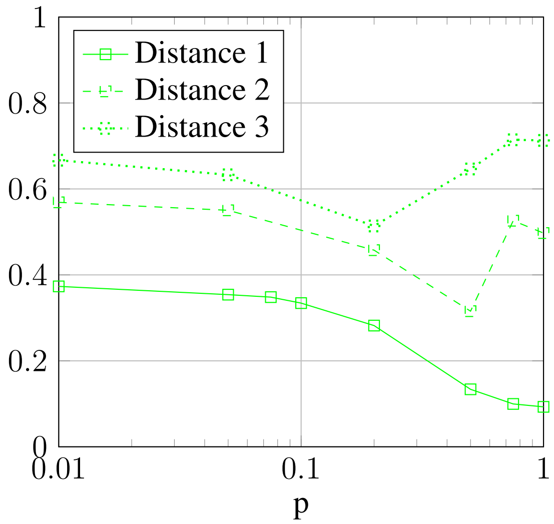

(c) Homogenisation

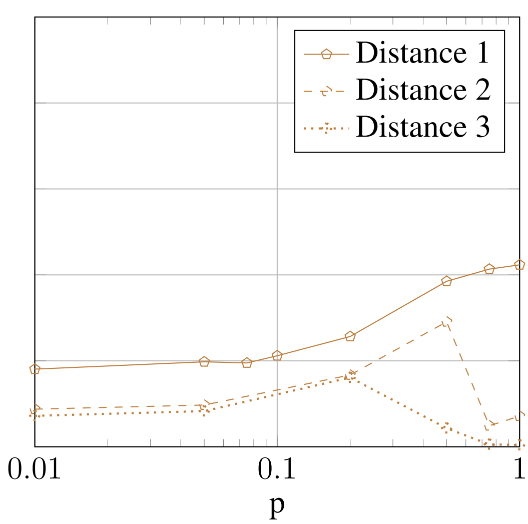

(d) Mixing

Figure 6.7: Results of path-based distance in Watts-Strogatz small-world networks ($degree = 12$**): When run on networks of progressively more rewired edges, initially the results of influence over distance follow the pattern of those where** $distance = 1$**: a decline in homogenisation as edges are randomly rewired. However, beyond around** $p = 0.1$**, the results start to diverge and influence over higher distance causes greater convergence (Fig. 6.7c) and less mixing (Fig. 6.7d). The distributions of results where** $distance = 2$ **are charted in Figure 6.8.**

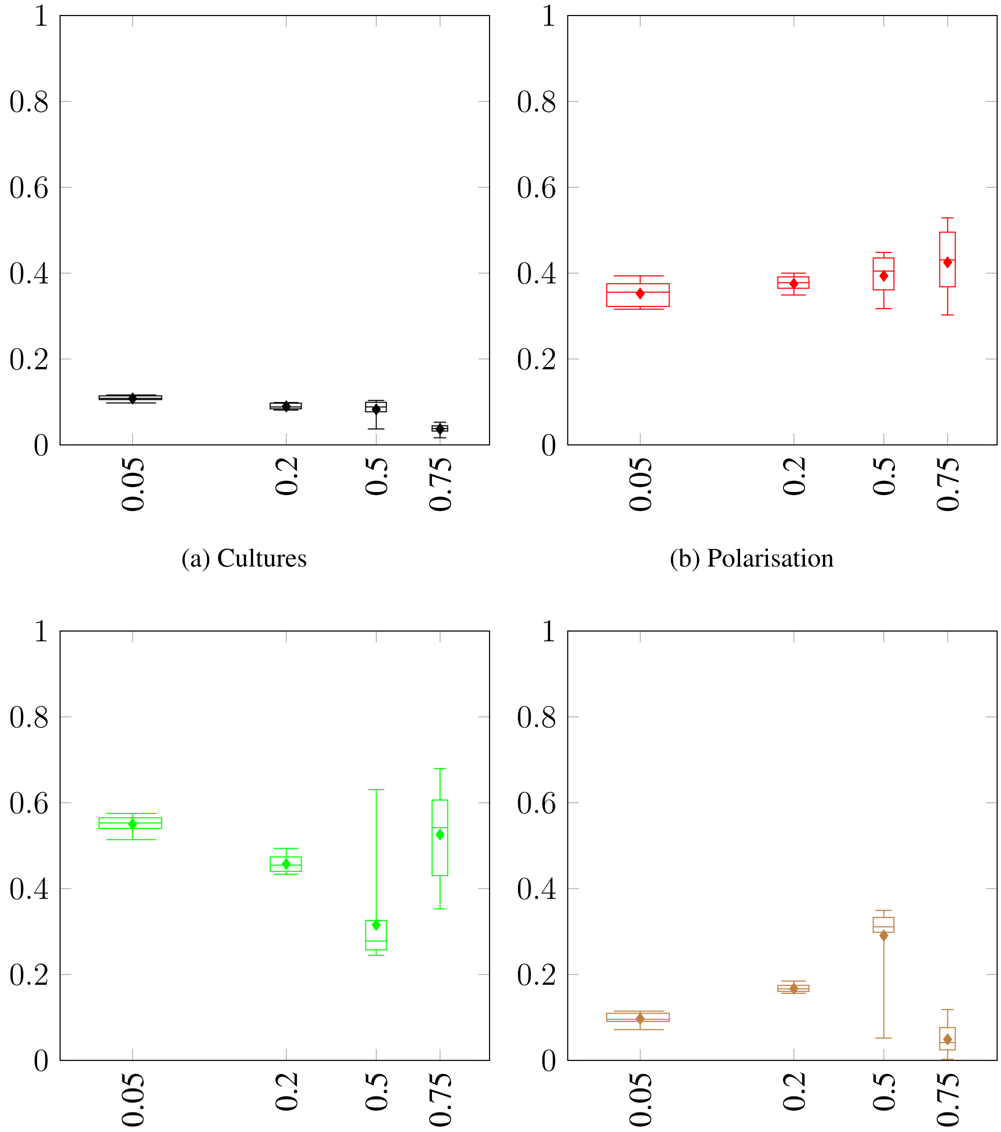

(a) Cultures

(b) Polarisation

(c) Homogenisation

(d) Mixing

Figure 6.8: Interquartile ranges for results of path-based distance in Watts-Strogatz small-world networks ($degree = 12, distance = 2$): Distribution of data given in Figure 6.7, for $distance = 2$. Despite a greater range of values for $p = 0.5$ and $p = 0.75$, the medians and interquartile ranges still show a clear 'u-shaped' pattern for homogenisation as the graph is rewired. The extended maximum value where $p = 0.5$ is the result of a single outlier.

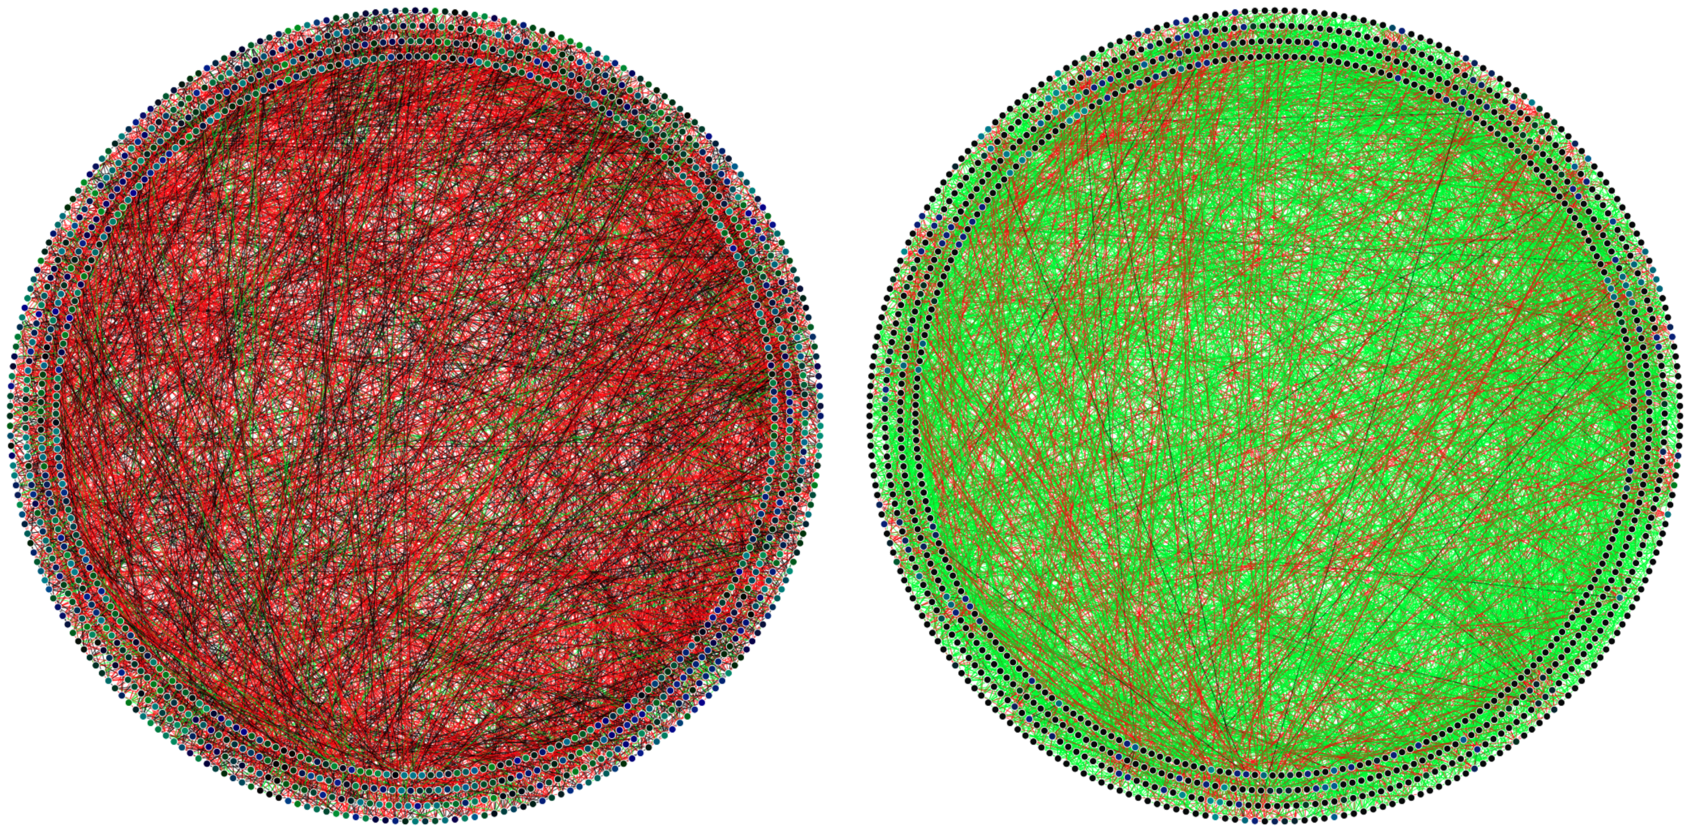

(a) Distance 1

(b) Distance 3

Figure 6.9: Visual overview of network showing differences when using path-based distance: Both networks depicted are Watts-Strogatz networks of degree 12, where each edge is rewired from the outer ring lattice with probability 0.5. The image on the left, (a), depicts a typical end state of the simulation where influence is only over distance 1 - i.e. direct connections. The image on the right, (b), depicts a typical end state of the simulation where influence is over distance 3. Polarised edges are red, homogenised green, mixed black. Where $distance = 1$, the rewiring of edges across the lattice has reduced homogenisation. Where $distance = 3$ there is far greater convergence of cultures: the number of cultures shrinks to 26, the largest of which comprises over half of the population.

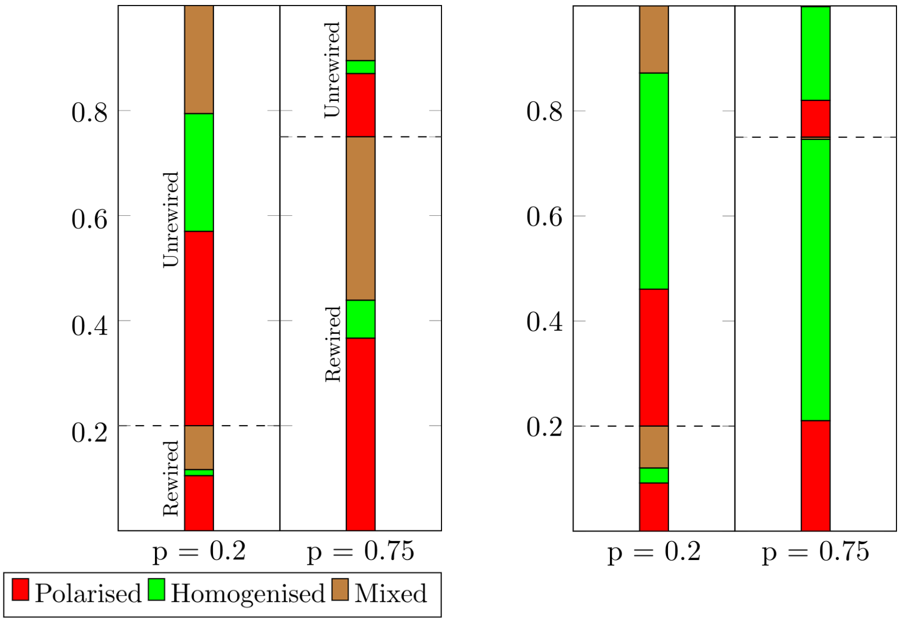

(a) Distance = 1

(b) Distance = 3

Figure 6.10: Composition of rewired and non-rewired edges using path-based distance for $distance = 1$ **(a) and** $distance = 3$ **(b). Rewired edges are represented below dashed line, non-rewired edges (i.e. those that remain part of the ring lattice) above dashed line.** p **is the probability of an edge in the lattice being randomly rewired. Where** $distance = 1$**, homogenised links decline with the reduction of unrewired edges in the ring lattice (6.10a right). Where** $distance = 3$**, as more of the lattice is rewired, mixing all but disappears for both rewired and non-rewired edges (6.10b right).**

rewired edges that link homogeneous cultures explodes, while homogeneity remains at a similar level for non-rewired edges of the ring lattice. The increase in both the proportion of rewired homogeneous links, and the increase in the total number of these rewired links, contribute to an increase in homogenisation and cultural similarity at the expense of cultural mixing (Figure 6.10b).

D: The effect of inverse-square-based distance in Watts-Strogatz small-world networks

When we apply an inverse-square based distance mechanism, we also see an increase in homogeneity with rewiring (Figure 6.11). However, when compared to the path-based-distance method, this increase occurs at a slightly lower level of rewiring (p = 0.075 for $degree$ = 12, d = 3). At higher levels of rewiring, more extreme changes may occur. While in path-based-distance, cultural mixing may recover very slightly at around p = 1; when using an inverse-square process cultural diversity not only increases at high levels of rewiring but becomes arguably the dominant dynamic. For $degree$ = 12 and distance d = 3, where rewiring probability p = 0.2, mixed links are barely above zero (Figure 6.11d); the average number of unique cultures is merely 23 in a population of 1024 agents. Yet at p = 0.5 (i.e. 50% of ring lattice links are rewired across the ring), the average number of unique cultures jumps to 1018 - almost 100% (Figure 6.11a), and the amount of 'mixed' edges is around 40% of the links in the network. Homogeneous edges between completely similar nodes all but disappear. The number of successful copying interactions which take place using such a parameter set is very low: agents tend to retain their starting traits. This abrupt change appears similar to a phase transition. Furthermore, at the transition point p = 0.2, the cumulative distribution of culture sizes appears to be close to a power-law.

6.4.3 Discussion

The results from distance-based influence on ring lattices suggest subtle differences between *inverse-square* and *path-based* distance methods (sections 6.4.2 A, 6.4.2 B). While both saw increases in homogeneity, this was higher for path-based distance, and increased even further with greater node-degree. One probable cause is a fundamental difference between how influence-over-distance is determined: with inverse-square, each node within d is considered once; with path-based-distance, each *path* of length

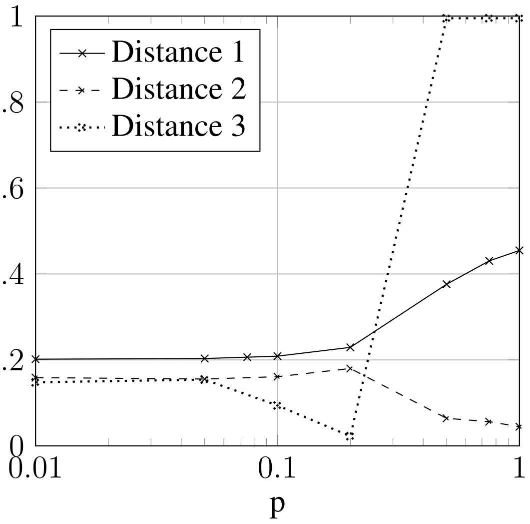

(a) Cultures

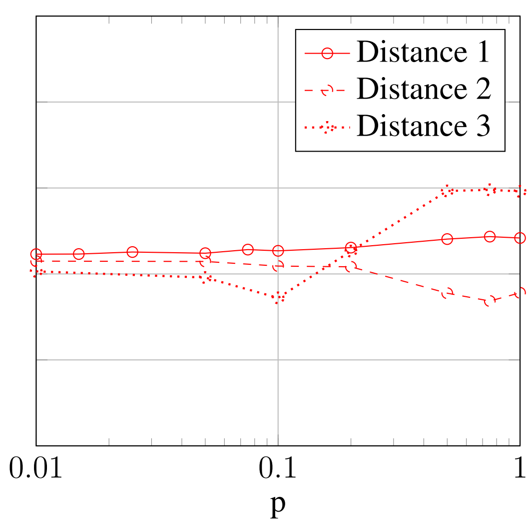

(b) Polarisation

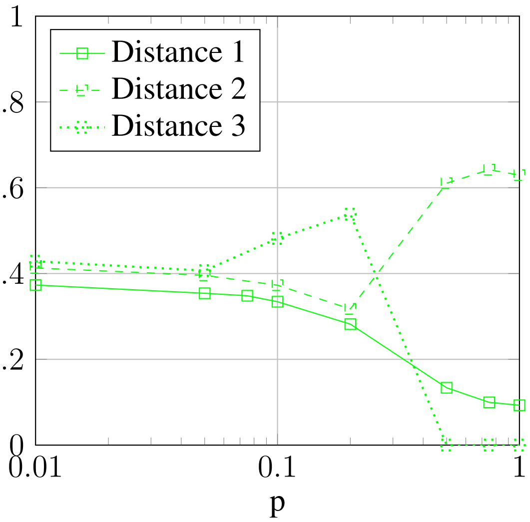

(c) Homogenisation

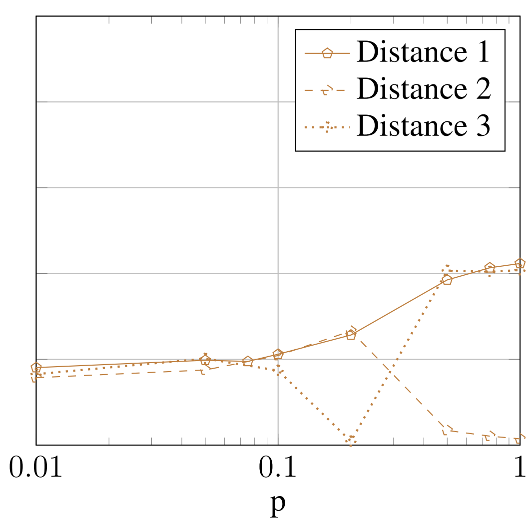

(d) Mixing

Figure 6.11: Results of inverse-square-based distance in Watts-Strogatz small-world networks ($degree = 12$): Initially, results follow a similar pattern to path-based distance (Figure 6.7), where increasing rewiring and distance causes greater convergence (Figure 6.7c). At $p = 0.2, d = 3$, the number of unique cultures has shrunk to a very small number and the amount of mixing is low: most agents are identical to their neighbours or completely dissimilar. However, with only a little more rewiring there is then a jump to an almost opposite state: at $p = 0.5, d = 3$ the population is culturally fragmented and agents tend to retain their traits.

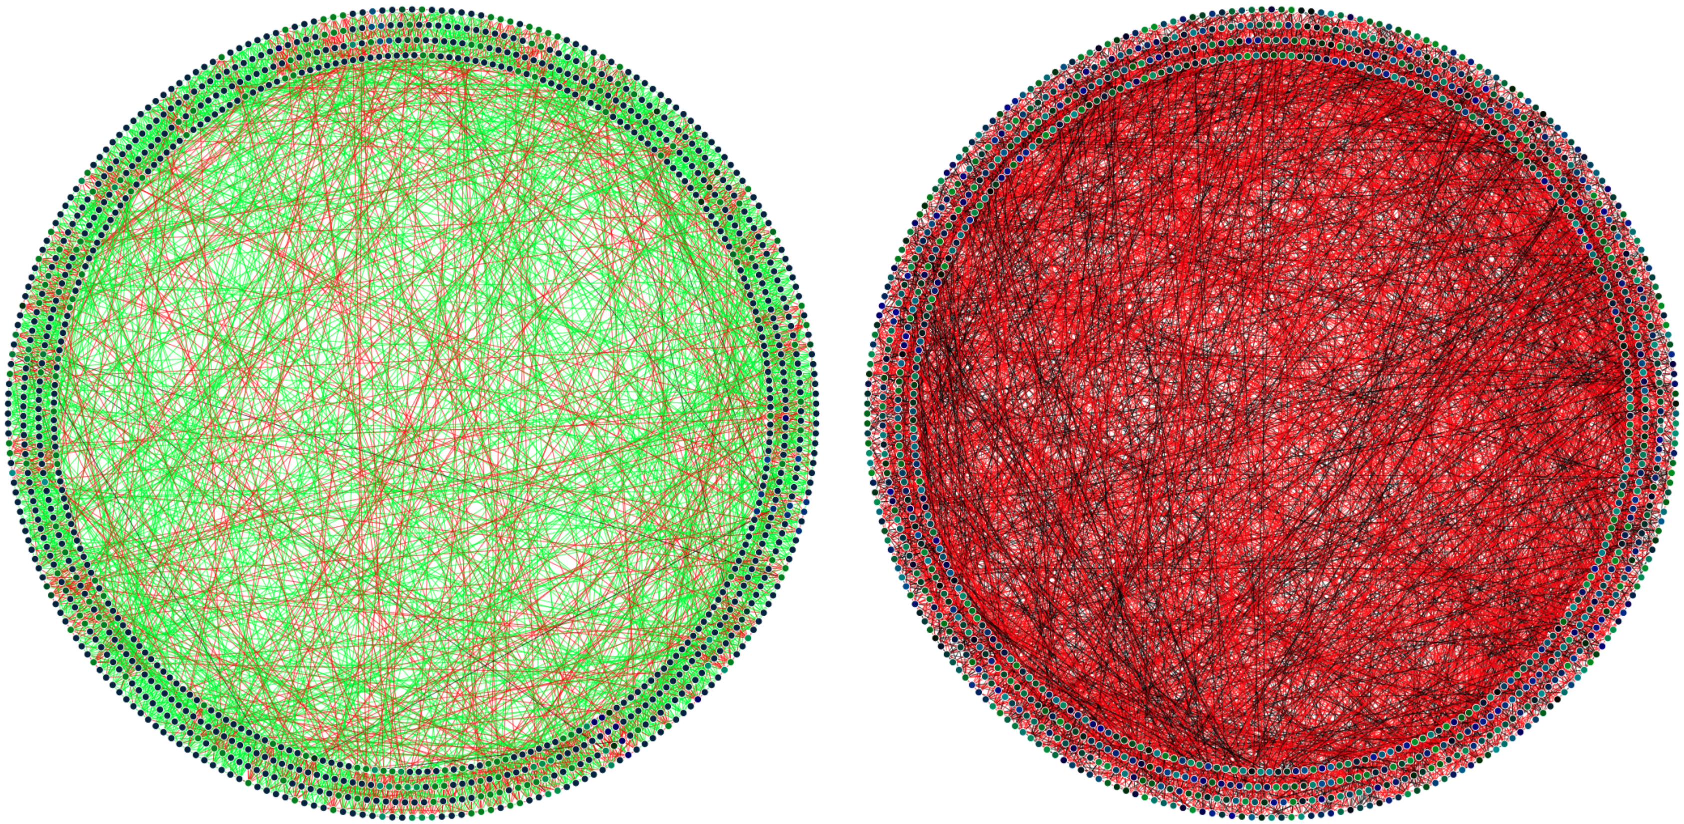

(a) Rewiring probability $p = 0.2$

(b) Rewiring probability $p = 0.5$

Figure 6.12: Visual overview of network illustrating transition from homogeneity to cultural fragmentation when using inverse-square distance: Depicted are two Watts-Strogatz networks of degree 12, stabilised at the end of our simulation using inverse-square-based distance of $distance = 3$. On the left, (a), rewiring edges from the ring lattice with probability $p = 0.2$ has caused a large degree of homogenisation (green edges) and convergence. Here, there are few unique cultures (22) and the size of the largest culture is over half of the population. However, with a little more rewiring $p = 0.5$ (right), the simulation instead stabilises in a state of cultural fragmentation: with 1019 unique cultures, each agent's culture is almost unique from any others. Agents may have some traits in common, but rarely, if ever, complete similarity.

$\leq d$ is considered. Many of these paths will be to the same nodes multiple times, and even more so to the nearest neighbours who have acquaintances in common. This is likely to reinforce and advance popular traits within local clusters.

When rings are rewired, the results where influence occurs at distance are non-monotonic, and less predictable than those where $distance = 1$. The dynamics described above, which tend to increase homogenisation in ring lattices, are also at play in rewired small-worlds. When influence is only via direct connections, rewiring (and thus reducing path length and clustering) tends to increase mixing - as seen in Chapter 4. However, when we introduce distance, each rewired link is not just allowing influence from its opposite vertex, but also that vertex's neighbourhood and beyond, driving homogenisation.

This increase in homogenisation seems to occur beyond the 'small-world window'[2], only when clustering has fallen significantly. In addition to sensitivity to clustering (Chapter 4), results could be due to already shortened paths becoming *wider* with additional rewiring. In small-world Watt-Strogatz networks, the length of the average shortest path drops significantly with only small amounts of rewiring; further rewiring then has a negligible effect on this measure [160]. This is most pertinent to traits which spread via simple contagion, where only a single path is needed [38]. However, as our model is more affected by multiple sources, additional rewired links across the lattice may provide the extra influence needed to tip an agent toward a culture.

The abrupt jump in 6.4.2 D (figures 6.11, 6.12) from a small number of cultures to almost complete cultural fragmentation appears similar to a phase transition, commonly observed in social models. The power-law distribution of culture size at the point of transition is also reminiscent of other studies based on Axelrod's model which exhibit phase-transition-like behaviour [35]. As we have seen with ring lattices, differences in results can be caused by inverse-square-distance adding influencing *nodes* over distance in contrast with path-based-distance adding *paths*. In a pure ring-lattice,

[2]The small world measure ω [151] is in the small world range of $-0.5 < \omega < 0.5$ for around $0.05 \leq p \leq 0.2$ where degree is 8-12, 1024 vertices.

increasing the distance of influence using inverse-square only adds a small number of additional nodes at each distance, as reach of influence inches around the ring: the number of influencing nodes is simply $distance \times degree$. When edges are rewired however, the number of additional influences increases non-linearly; although not quite exponentially as there is some redundancy in the rewiring of edges to the same local clusters. In such cases the number of successful influencers when using the inverse-square method can start to outnumber those when using *path-based-distance*, as in the latter distant influences can 'blocked' by dissimilar agents along the path. So in this scenario of high inverse-square-distance and high rewiring, the sheer number of mostly unique influencers swamps the activated agent; competing traits tend to cancel each other out. The number of traits being $q = 10$, and features being $F = 5$, means the probability of two agents starting with more than two feature/traits in common is close to zero. Where the number of agents is 1024, and p is high enough that an agent has links to a large proportion of others, the agents most likely to influence an agent on a feature are those *which already hold the same trait*. Thus agents largely retain their starting traits, and no culture exerts enough influence to change them. There is a tipping point from where agents have few enough influences to successfully copy from the most influential, to where influence is diluted enough to inhibit action.

6.5 Conclusion

In this chapter we have extended our model, adding the concept of influence over distance. We introduced two possible methods of incorporating influence beyond directly connected nodes, and repeated the experiments of Chapter 4 with these included. How these often weaker, but often more numerous, additional influences affect cultural spread is dependent on the network structure.

When clustering is high and contacts are predominantly local, strong social norms may form. However, individuals who seem relatively dissimilar to their local group may be alienated; invoking Simmel's stranger appearing both *near* (in terms of proximity) and *far* (in terms of similarities) to his social group. There are existing examples of increased contact either causing convergence of traits [6, 132], or conversely, further entrenching existing positions [8, 10]. Our simulations model scenarios where as individuals expand their horizons and connect to others further away, they may become either more like their contacts or maintain their differences (Figures 6.7, 6.11). In [159], Duncan Watts invokes Asimov's planet *Solaria* to describe a world where individuals interact almost exclusively at distance and local clusters do not exist - an extremely 'rewired' scenario. Such an extreme example of predominantly long and weak ties in our simulations produces results where influence is so thinly spread that agents usually maintain their original traits; for each feature there are enough influencing agents to reinforce the current trait, but not enough similarity to induce homophilic copying (Figure 6.12b). 'Cultures' are almost entirely individualistic, dominant social norms do not seem to exist. Although such a scenario is perhaps unrealistic, there are nevertheless implications of connectivity in the consideration of the spread of phenomena such as misinformation: it may be easy to find connections to those who confirm rather than challenge our biases. Furthermore, the transition from a population of a small number of homogeneous cultures to one of high cultural fragmentation appears to occur as a phase transition, a phenomenon observed in other social models.

Finally, it is notable that with the exception of Nowak et al. [122], the notion of additional weaker influence over distance appears to be rarely included in social models. It may be the case that similar mechanics could reveal interesting behaviour in other social simulations.

Chapter 7

An empirical test of our model on real-world data

In previous chapters, our model has been used to examine abstract scenarios. To some extent, this is an intention of such models; positing that micro-specifications may be potential explanations of macro-behaviour. Axelrod's culture model showed that global polarisation and the persistence of different cultures can exist even where local rules promote convergence [6]. In this sense, the model could give indications as to the probable number of cultures at stabilisation, but it was likely not intended to predict which cultural traits would succeed, or localised results such as the final state of any individual agent. The micro-level rules were the focus, rather than the trait values themselves, which were allocated randomly at initialisation of the simulation. This is a common approach of agent-based models, to isolate specific (and often simple) local behaviours as candidate explanations for global behaviour [57].

However, the use of agent-based models as social simulations need not be confined to abstract or qualitative studies. There are examples of the approach applied in a more empirical context, on real-world data. When modelling stock markets, Bak et al [12] found that allowing agents to imitate each other causes price distributions consistent with the real-world data, while Kirman and Vriend modelled price dispersion and buyer loyalty in a fish market [83]. Axelrod and Bennett [7] recreate history in the modelling of European alliances in the 1930s, and also offer an alternative counterfactual history where allied nations instead unite against the Soviet Union. Axtell et al [9] also constructed an agent-based model that was compared with historical data; one which retrospectively predicted the population decline of the Native American Anasazi culture. The use of such a model may offer some explanations as to the factors affecting such a decline.

In this chapter, we aim to test our model in a similar fashion: *given starting conditions, we ask to what extent will our model arrive at a similar later state to that of the real-world data.* This chapter supports research question RQ7.

7.1 Selection of a suitable dataset

If we are to compare results generated by our models against real world data, then we need both starting conditions and data of the same structure gathered at a later date. In this way, we can take the data at an early time-step, run them within our model and compare with real-world data at a later time-step; hopefully, the results from our models will show some similarity with how things actually occurred. Therefore, any dataset must have a temporal, longitudinal aspect.

As our work principally explores agent influences on each other, and the wider effects within networks, the data must be able to form a network graph of agents, each node having stated or derived connections with a subset of the population.

Thirdly, these individual nodes or agents must hold multiple mutable attributes that may be changed over the course of the longitudinal study. The nature of Axelrod's culture model and its derivatives is such that each changeable feature should have possible values (traits) that are categorical, non-ordinal, and ideally mutually exclusive. The latter attribute can be harder to discern; for example most football fans may have a favourite club but some may support two, individuals may hold multiple nationalities. In such instances it may be possible to model these cases in terms of an entity's most strongly held preference.

It has proven difficult to find data that exhibits all three of these characteristics, an observation also acknowledged by Lewis et al. in their creation of the 'Tastes, ties, and time' (T3) dataset [93]. The title of such a dataset seems to identify it as ideal for our purposes; indeed the authors define the dataset as sociocentric, multiplex, and longitudinal. However, the methods of data collection and anonymisation attracted public criticism, and the dataset was withdrawn [165]. It may be possible to 'scrape' online social media to construct networks of individuals with stated beliefs and preferences, but clearly this is ethically problematic [165].

The most suitable dataset we have yet found is the 'Glasgow Teenage Friends and Lifestyle' dataset [145], which we shall use as an empirical test of social influence models.

7.1.1 The 'Glasgow Teenage Friends and Lifestyle' dataset

The Chief Scientist's Office of the Scottish Home and Health Department, under the remit of their Smoking Initiative, funded a study of the lifestyle and friendship groups of teenagers in the West of Scotland in the 1990s [105, 31, 129, 130]. Questionnaires given to teenagers recorded their use of alcohol, tobacco, and drugs, and also their leisure activities and friends. An excerpt of this data forms the basis for the 'Glasgow Teenage Friends and Lifestyle' dataset [145].

This dataset focuses on a cohort of pupils at an unnamed comprehensive school in the wider Glasgow area. The study took place from February 1995, when the students were aged around 13, to January 1997. Questionnaire responses were recorded at three intervals, giving longitudinal data regarding how behaviours may have changed over the time of the study. The dataset includes, for each time-step, an adjacency matrix representing the friendship network of the participants. Thus we have temporal data, consisting of a connected network of nodes with mutable attributes.

For a list of fields included, see Table B.1 in Appendix B.

7.2 Methodology

We extract features and network edges from the 'Glasgow Teenage Friends and Lifestyle' dataset. Those present at the first time-step will form the starting 'seed' data for initial configuration of our models. The extraction of these features and edges is discussed in more detail in Section 7.3.

When the simulations have completed, we will compare the resulting feature vectors for each agent with those of the third and final time-step of the real dataset. We define $consistency$ as the degree to which a simulation matches the real data. Where all features for an individual in the simulation results match all features for the same individual in the real data, this would be $consistency = 1$. We run our simulations multiple times using different random seeds, and compare the results against the real data each time, before taking an average of the consistencies. Note that the random seed is only used to determine the order in which the agents are activated within the simulation.

We test three models against the data; Axelrod's original culture model [6], an adapted version of Flache and Macy's social influence model (see Section 3.3.4) [59], and our social influence model incorporating multiple influence and influence over distance. For our social influence model, we use distances 1 to 3.

7.3 Data preparation

The 'Glasgow Teenage Friends and Lifestyle' dataset is provided by the SIENA project, and consists of several files in the R Data format [145]. The data were extracted from these files, cleaned and saved in Graph Modelling Language (GML) for use by our model.

We focused on those students included in *selection129*, i.e. those who were present at all three time-steps. Others were excluded.

7.3.1 Creation of features and traits

When selecting nodal attributes, those that are discrete and mutable (specifically 'influenceable') are most appropriate for our models of social influence. We excluded all geographic data; while potentially useful in determining the likelihood of links forming (e.g. pupils living close by forming friendships), we considered it unlikely to affect trait adoption beyond the friendships already specified. While pocket money could conceivably be influenced by the social norm of a pupil's friendship group, it is both a continuous value and one arguably determined more by a pupil's parents; therefore this too was excluded. The dataset records a pupil's age, sex, and the smoking habits of their family as immutable (or at least, not easily influence-able by their peers during the study).

For many fields (see Table B.1), values are given as a frequency with which an activity is undertaken. For example, for each leisure activity, the recorded responses were 'most days', 'once a week', 'once a month', 'less often or never'. As models based on the Axelrod culture model work on discrete and categorical traits, we cannot treat these values as ordinal. While 'most days' is closer to 'once a week' than 'once a month', we must instead treat these values as categorical and unrelated. As such, we have reduced these responses to binary values: either the pupil engages in the activity fairly regularly, or not.

We must also determine what constitutes 'regularly' engaging in an activity. Some activities, for reasons of cost and logistics, do not take place every day - such as pop concerts. Therefore we considered that attending a gig or pop concert at least once a month would indicate it is a favoured activity for that pupil. Conversely, due to the greater ease of listening to CDs we consider it a regular activity if a pupil undertook this activity at least once a week. When assigning these thresholds we used our own discretion informed by the distribution of pupil responses for each field. A list of these thresholds is given in Table 7.1.

Attribute	**Values**	**Considered TRUE where:**
Tobacco	1 (non), 2 (occasional), 3 (regular)	≥ 2
Alcohol	1 (non), 2 (once or twice a year), 3 (once a month), 4 (once a week), 5 (more than once a week)	≥ 3
Cannabis	1 (non), 2 (tried once), 3 (occasional), 4 (regular)	≥ 3
I listen to tapes or CDs	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I look around in the shops	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I read comics, mags or books	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I go to sport matches	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 3
I take part in sports	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I hang round in the streets	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I play computer games	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I spend time on my hobby (eg art, an instrument)	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2

Attribute	**Values**	**Considered TRUE where:**
I go to something like B.B., Guides or Scouts	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I go to cinema	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 3
I go to pop concerts, gigs	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 3
I go to church, mosque or temple	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 3
I look after a pet animal	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 2
I go to dance clubs or raves	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 3
I do nothing much (am bored)	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	≤ 3

Table 7.1: Thresholds for converting data to binary trait values

We created four feature-sets for experimentation, each described below. We acknowledge that models based on Axelrod's culture model typically have several possible trait values (q) rather than just binary data. Therefore we have attempted to increase q beyond 2 by grouping music genres and types of leisure activities.

Feature-set 1: All mutable fields

This feature-set consisted of a feature for each mutable field, each with a binary value. The assumptions on what constitutes a true or false value are given in Table 7.1. Therefore this feature-set is of $F = 34, q = 2$.

Feature-set 2: A selection of four fields

A simplified version with a small number of fields: tobacco smoking, cannabis use, alcohol use and participation in sports. This is based on the 's50' excerpt of the full dataset [144], but with the full friends network retained. All trait values are binary, giving $F = 4, q = 2$.

Feature-set 3: All mutable fields, with grouped music genres

To increase q beyond 2, we group music genres into 6 trait values using criteria based on work by Rentfrow and Gosling [135]. Musical preferences were grouped into categories: *'Intense and Rebellious'* (heavy metal, indie, rock, grunge), *'Reflective and Complex'* (jazz, classical, folk/traditional), *'Upbeat and Conventional'* (chart music), and *'Energetic and Rhythmic'* (reggae, rap, hip-hop, techno, house, rave, dance). Where a pupil's responses indicated a preference for one of these categories, this was recorded as the trait value. In some instances several categories scored equally, and were given the music trait *'Eclectic'*. Also, a small number of students did not claim to listen to any music genre, and were given the trait *'Nothing'*.

Musical category	**Dataset time-step 1 total**	**Dataset time-step 3 total**
Intense and Rebellious	14	23
Reflective and Complex	2	0
Upbeat and Conventional	41	29

Musical category	**Dataset time-step 1 total**	**Dataset time-step 3 total**
Energetic and Rhythmic	21	17
Eclectic	50	60
Nothing	1	0

Table 7.2: Distribution of preferred musical category over 129 students.

Other mutable fields were used in the same binary way as in other feature-sets. This gives $F = 20, q = 6$ (although for 19 features, only two trait values will be used).

Feature-set 4: All mutable fields, with grouped music genres and grouped leisure activities

In addition to grouping musical genres, we also attempt to group types of leisure activity using categories based on work by Yin et al [163] and Agnew and Petersen [1]. Leisure activities were grouped as *'Unsupervised socialisation'*, *'Organised leisure activities'*, *'Organised sport'* and *'Self-directed'*. Where students showed an equal preference for several activity types, the trait was recorded as *'Eclectic'*. How the various leisure activities were categorised is detailed in Table 7.3.

This gives us $F = 6, q = 6$ where we have the music field grouped into 6 genres as above, an activities field grouped into 5 possible values, and four binary fields: alcohol use, cannabis use, tobacco use and romantic.

Activity	**Category**
I listen to tapes or CDs	Self-directed
I look around in the shops	Unsupervised socialisation
I read comics, mags or books	Self-directed

Activity	Category
I go to sport matches	Unsupervised socialisation
I take part in sports	Organised sport
I hang round in the streets	Unsupervised socialisation
I play computer games	Self-directed
I spend time on my hobby (eg art, an instrument)	Organised leisure activities
I go to something like B.B., Guides or Scouts	Organised leisure activities
I go to cinema	Unsupervised socialisation
I go to pop concerts, gigs	Unsupervised socialisation
I go to church, mosque or temple	Organised leisure activities
I look after a pet animal	Self-directed
I go to dance clubs or raves	Unsupervised socialisation

Table 7.3: Grouping of leisure activities into categories.

7.3.2 Creation of edges

For each feature-set, edges between nodes were created in two ways; undirected and unweighted, directed and weighted. Once our simulation has been run on the starting conditions obtained from time-step 1, we will use the resulting trait values as starting traits for time-step 2, albeit with the friendship edges given in the adjacency matrix for time-step 2. In other words we will run our simulation in steps, with friendship edges changing for each step as they do in the dataset. Where we use the term ‘time-step’ in this chapter, we are referring to time-steps in the real data; we refer to sequential changes in a run of our simulation as ‘iterations’.

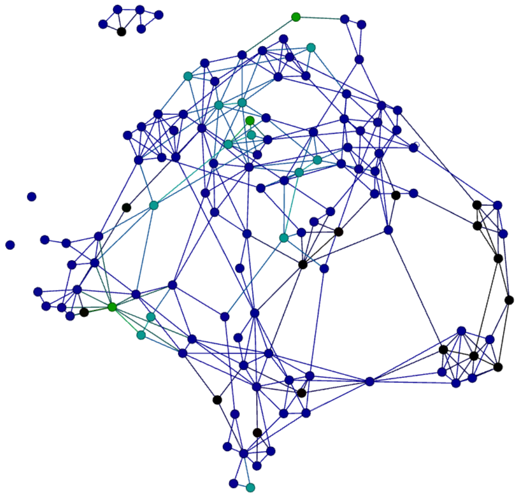

Figure 7.1: Undirected edges derived from the Glasgow Teenage Friends dataset.

Undirected and unweighted

Using the friendship adjacency matrices, where a friendship existed between two pupils in any direction, a reciprocal undirected edge was created. These edges were unweighted, regardless of whether pupil was identified as 'friend' or 'best friend'. This network is illustrated in Figure 7.1.

Directed and weighted

Using the friendship adjacency matrix for time-step 1, directed and weighted edges were created. Where a pupil described another as a 'friend', an out-edge of weight 1 was created. Where a pupil described another as a 'best friend', an out-edge of weight 2 was created. i.e. a pupil will be influenced only by those they have described as 'friends' or 'best friends'.

7.4 Results

7.4.1 Undirected and unweighted edges

All mutable, binary traits; undirected edges

Treating all mutable fields as binary value traits ($F = 34, q = 2$), the probability of any randomly assigned trait matching a corresponding trait in the dataset is simply 0.5. Where edges are undirected and given equal weight, all models perform better than random chance (Figure 7.2a), with the models of multiple social influences - ours and Flache and Macy's - obtaining greatest consistency with the real data.

Small selection of fields; undirected edges

When a smaller selection of just four fields[1] is used, any differences in results from model to real-life are magnified. This is reflected in the greater variance of results and lower worse-case performance (Figure 7.2b). Despite this, median values remain comparable with the results of all mutable ($F = 34$) traits.

7.4.2 Directed and weighted edges

It should be noted that when run on a directed network, the Axelrod model, and in most cases the Flache and Macy model[2], did not stabilise within a reasonable amount of time. For the former we allowed 10^9 iterations, and the latter 10^7. In fact, the Axelrod model (and likely Flache and Macy also) would never stabilise on the directed network created by the dataset values at time-step one. The reason for this is the existence of a small number of nodes with in-edges but no out-edges. These agents

[1] Alcohol use, smoking, cannabis use, and participation in sports

[2] This is despite the generous stop condition which we described in Section 3.3.4, even where this was set to allow 10,000 dormant iterations per agent.

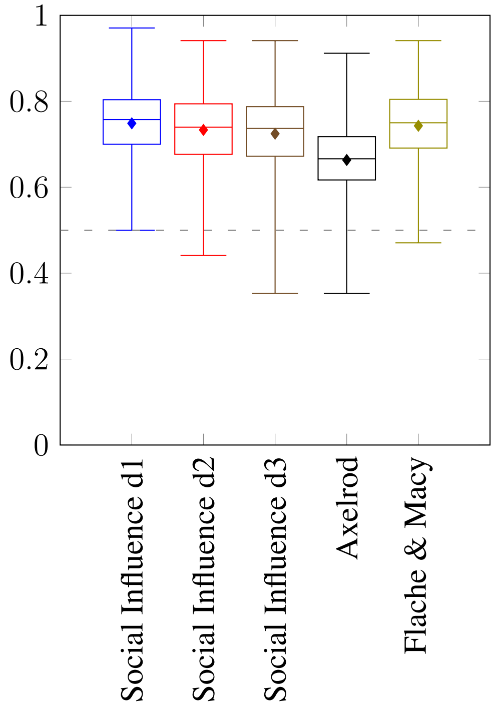

(a) All mutable, binary traits

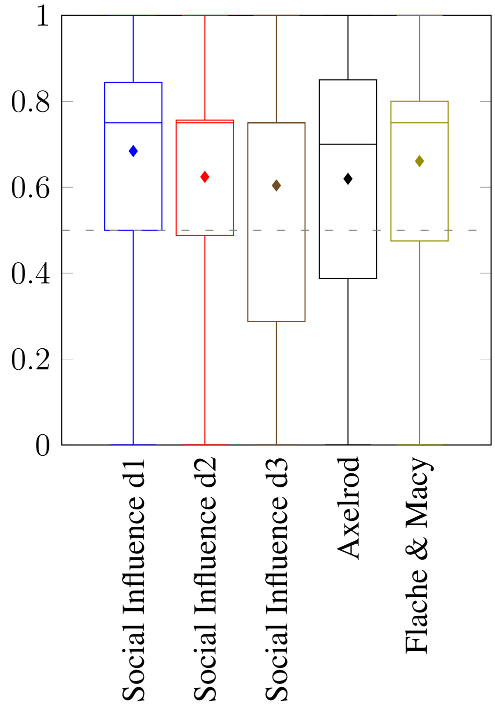

(b) Selection of four fields

Figure 7.2: Consistency between model and real data: undirected edges. Presented are the similarities between the end states of simulations, and the actual traits in the dataset at time-step 3. Dashed line indicates the probability of randomly assigned traits matching those of the dataset (i.e. in this case, binary traits for all features results in 0.5). All models on average perform better than random chance, with the better results from our social influence model where influence-distance = 1.

are copied *from* but never copy any other agent themselves; they are influenc*ers* but never influenc*ees*. These uninfluenceable individuals have a degree of dissimilarity, essentially injecting competing traits into the wider population. These traits may move in waves throughout the population, but never reach fixation as agents continue to copy competing traits from the alternative, unchanging, uninfluenceable agents.

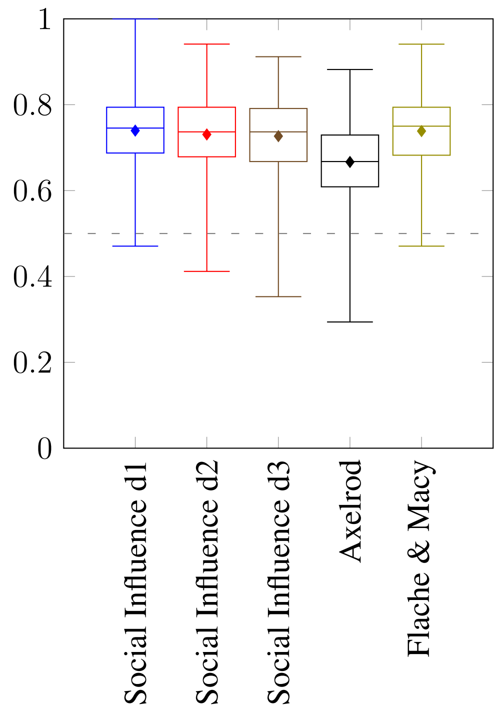

(a) All mutable, binary traits

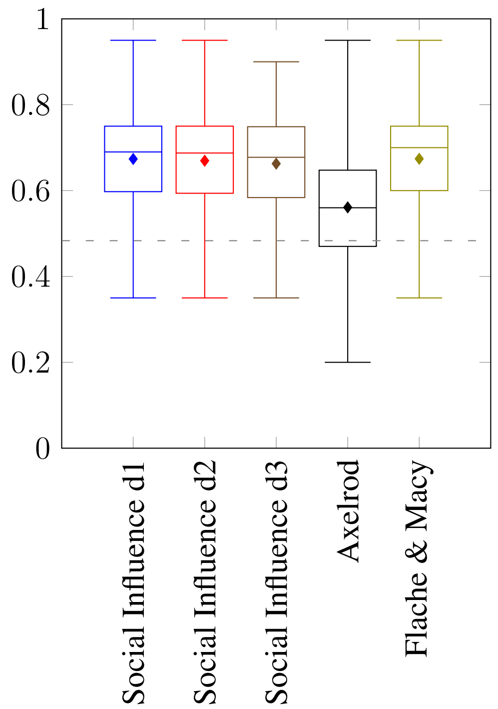

(b) Grouped music genres

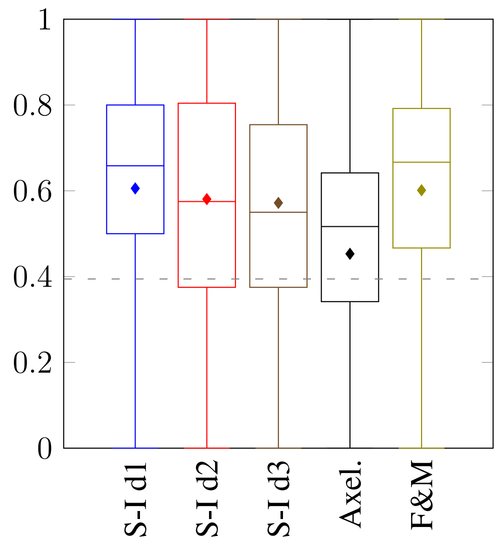

(c) Grouped music genres and grouped activities

Figure 7.3: Consistency between model and real data: directed edges. In most cases, models based on multiple sources of influence perform better than Axelrod's. A smaller number of fields intuitively increases variance (7.3b and 7.3c). Note that the probability of randomly assigned traits matching those of the dataset (dashed line), is in the cases of 7.3b and 7.3c lower than 0.5 as trait values are no longer binary ($q > 2$).

All mutable, binary traits; directed and weighted edges

The results for $F = 34, q = 2$ on a directed and weighted network are very similar to those for the same feature-set on an undirected network. Means, medians and interquartile ranges are almost identical. There is a greater distance between the minimum and maximum values (Figure 7.3a).

Grouped music genres, directed and weighted

The reduction in total features due to the grouping of 14 music genres into a single field causes any differences between feature similarities to have a proportionally greater effect. However for this feature-set, models of multiple influence show only a mild reduction in mean/median similarity (around 0.05). Variance increases for all models (Figure 7.3b).

Grouped music and activities, directed and weighted

The additional decrease in fields further lowers the ability of the models to match the real-life data. However, compared with the feature-set using 4 binary fields on an undirected network, the majority of results from models of multiple influence stay above the results of random chance (Figure 7.3c). Variance is significantly higher than when only grouping musical genres however.

7.4.3 Composition of edges

In previous chapters we have examined the proportion of edge states based on the similarity of agents they connect; the states being *homogenised*, *polarised* and *mixed* (see Section 3.3.2). When examining the real data for these edge-types, we find that for most of the above feature-sets, polarised edges are rare or non-existent at time-step one. At time-step three in the real data, they are even rarer. The vast majority of edges in the dataset are mixed, i.e. between agents with some but not all traits in common (see Table 7.4).

When run through the various models, polarisation remains low. However, convergence occurs changing some mixed links into homogenised. Factors which have been seen to drive homogenisation, such as increasing influence over distance (Chapter 6), also do so here. The greatest convergence is seen in Axelrod's model and our model where distance is 3 (Figure 7.4).

Feature set	**Dataset time-step**	**Edges**		
		Polarised	Homogenised	Mixed
Undirected, all mutable, binary	Dataset values, time-step 1	0	0.00154	0.99846
	Dataset values, time-step 3	0	0	1
Undirected, small feature selection, binary	Dataset values, time-step 1	0.012326656	0.362095532	0.625577812
	Dataset values, time-step 3	0.006349206	0.257142857	0.736507937
Directed, all mutable, binary	Dataset values, time-step 1	0	0.002227171	0.997772829
	Dataset values, time-step 3	0	0	1
Directed, all mutable, music grouped into genres	Dataset values, time-step 1	0	0.002227171	0.997772829
	Dataset values, time-step 3	0	0	1
Directed, all mutable, grouped music genres and grouped activity types	Dataset values, time-step 1	0.008908686	0.028953229	0.962138085
	Dataset values, time-step 3	0.004310345	0.051724138	0.943965517

Table 7.4: Edge composition of real-world dataset: The proportion of polarised, homogenised, and mixed links for the dataset at the first and last time-steps. Mixed links dominate; however there is a slight increase in convergence of the grouped music and activities feature-set.

7.5 Discussion

7.5.1 Multiple-influence models show greater consistency with real-world data

Results from all feature-sets show a greater consistency between the real-world data and multiple-influence models than they do to Axelrod's dyadic model. As we saw in Chapter 3, results from our model of multiple influence bear closer similarity to those from Flache and Macy's. Although both models feature influence from multiple sources, the processes in which copying of traits take place are very different; yet statistical measures are again comparable. In so much that our model may share characteristics with those of complex contagions (Chapter 4), these results may lend weight to the argument that the spread of behaviour often requires multiple sources of influence [36, 38].

The results of influence over greater distance in each case show lower accuracy than where $distance = 1$. This may, at least in part, be due to their tendency toward greater convergence and homogenisation; we have seen in Section 7.4.3 that the real data does not exhibit such a strong convergence. However, even in cases where influence at $distance = 3$ caused greater convergence than the Axelrod model (figures 7.4c, 7.4d, 7.4e), it still showed greater similarity with the dataset's final traits (Figure 7.3).

7.5.2 Instability of dyadic models when a minority are non-influenceable

In Section 7.4.2, we establish that in instances where agents are copied from, but do not copy from others, the Axelrod model may never reach a static state. There could be myriad social reasons, beyond the scope of this thesis, as to why these pupils were regarded as friends by others yet did not identify any friends themselves. They may

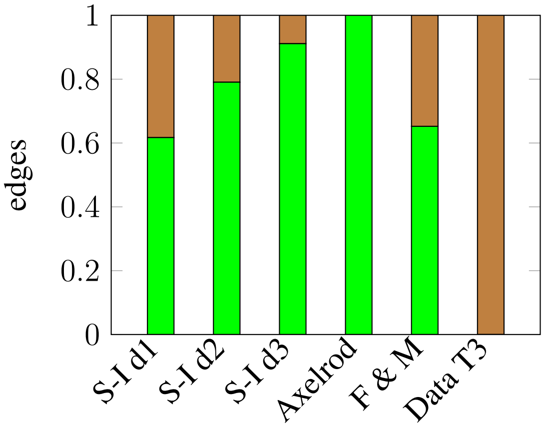

(a) Undirected, all mutable, binary

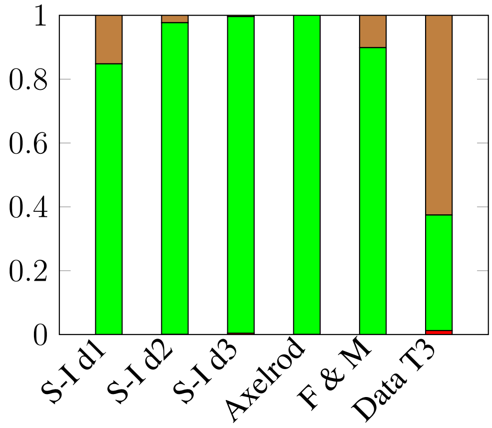

(b) Undirected, selection of 4 fields, binary

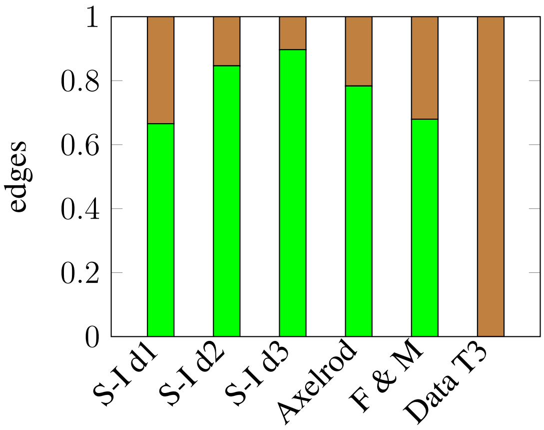

(c) Directed, all mutable, binary

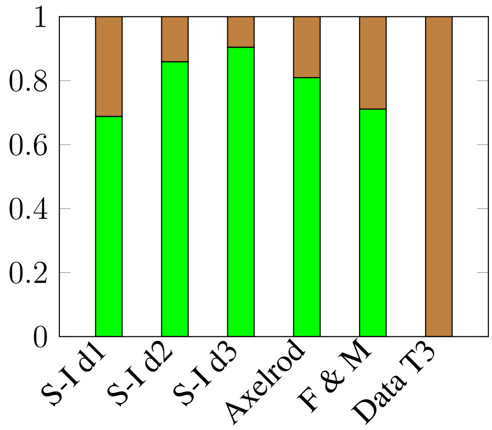

(d) Directed, grouped music genres

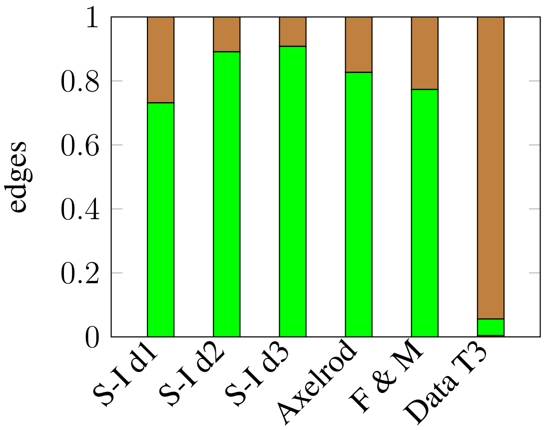

(e) Directed, grouped music genres and activities

Figure 7.4: Edge compositions for each model on each feature-set: The proportion of polarised, homogenised, and mixed links. Models illustrated are our social-impact model at distance 1 (S-I D1), distance 2 (S-I d2) and distance 3 (S-I d3); Axelrod's model, Flache and Macy's model (F & M). Also included is the composition of links in the dataset at time-step 3 (Data T3).

have thought themselves above stating a small number of friends, or it may simply have been a mistake of omission on the survey response. These results do however tie in with other works on ‘committed’ or ‘zealous’ [111] individuals or groups within a population. Singh et al [143] found that a small number of ‘committed’ identical and unchanging agents in a simulation of the Axelrod model caused a reduction in time to consensus. Other studies have shown that competing committed individuals or groups may prevent a stable majority, instead fluctuating between two or more states [111, 162]. Our results are closer to the latter; the different traits held by those un-influenceable individuals causes other agents to switch back-and-forth between them in waves throughout the wider population.

7.5.3 Polarisation is almost non-existent

When examining the real dataset, we found that polarised edges are rare (see Section 7.4.3). Given most feature-sets used have a large number of features, and relatively low number of traits, this is intuitive; the chances of two friends having nothing in common is very low. In fact, in the feature-set of 34 binary traits, the probability of two agents with randomly chosen traits having none in common is 5.8208×10^{-11}. The probability of two agents with 34 randomly chosen features having *everything* in common is, in the case of binary traits, the same figure. Yet the proportion of homogenised edges at time-step one is slightly higher, reflecting that linked individuals in this study have stated their friendship and are thus more likely to have the same interests. Even so, the vast majority ($> .96$ for most feature-sets) of connections are *mixed*, and at later time-steps in the data remain so; perhaps reflecting that even close friends and family will not agree on everything (Table 7.4). It seems that in most cases, the less a model tends toward (almost) total convergence and homogenisation, the closer its results will be to the real data. Even where the number of fields is low, the completely polarised and distinct cultures that characterise Axelrod’s model do not occur: different individual tastes and behaviours persist even when in contact with other ‘cultures’. As we have seen in chapters 3, 4 and 6, with our model cultural mixing usually continues to exist.

7.5.4 Limitations

The dataset did not exhibit a large amount of change over the course of the study. Similarity between traits at time-step 1 and time-step 3 in the real data was around 75%, dependent on the feature-set used. The performance of the models could in part be attributed to the Axelrod model causing many more changes between agents than ours. Nevertheless, at each run a large number of agent interactions comprising successful trait copying did take place in the simulations using our model.

Our attempt to raise the number of possible trait values (q) may have limitations. Any grouping into genres and categories requires some subjective discretion. There are a large number of different ways in which we could have grouped this data, all possibly giving different results - better or worse. We believe that those we have chosen are reasonably representative. Despite the grouping of these preferences reducing the number of fields and thus magnifying any differences, the models still scored higher on this feature-set (Figure 7.3c) than another low-F feature-set (Figure 7.2b) which didn't include grouped music genres and activities.

A flawed simplification of these groupings is that the music genres or activity types are treated as if they are almost mutually exclusive. We chose the field value based on a student's strongest preference, and included an 'eclectic' trait for those pupils with a broad range of interests. However, a small change in survey response could tip a student from 'eclectic' into another category. Some activities may also be complimentary; for example one who goes to gigs (classified as 'unsupervised social') is also likely to listen to CDs (classified as 'self directed'). This is one of the main difficulties we have found in applying extensions of Axelrod's culture model to real data. The model assumes multi-dimensional data of multiple, discrete, mutually exclusive values. On the other hand, the use of a continuous but often one-dimensional scale (as often found in opinion-formation models) is also a simplification. A refined model may manage to combine the strengths of both approaches, allowing homophilic copying of multi-dimensional ordinal data.

7.6 Conclusion

In this chapter we have taken models usually employed in more abstract scenarios, and applied them to longitudinal real-world data. Despite limitations, when given real data as a starting seed, in most cases these models do produce a good approximation of the later longitudinal data. Models incorporating multiple sources of influence performed better than those using only dyadic influence. However, applying influence over distance produced slightly degraded results. Interesting behaviour emerges when dyadic models are run on directed networks; 'committed' un-influenceable agents cause the simulation to persist in a non-static state. Our more deterministic model of compound influence did allow such a scenario to stabilise to a static state. When examining the real, longitudinal data, the completely polarised worlds famously generated by Axelrod's study did not materialise. Instead, agents moved more toward diversity: a key factor in the greater consistency between the data and the results of multiple-influence models than between the data and Axelrod's model.

Finally, it proved difficult to find datasets comprising suitable longitudinal, networked, multi-dimensional and mutable data. The 'Tastes, ties and time' study [93] recognised this, and it is hoped that researchers can solve the ethical and anonymisation issues associated in obtaining and presenting such data. If handled sensitively, its use in computational social science could be of great benefit.

Chapter 8

Conclusions

The central theme of this thesis has been the embedding of Social Impact Theory in a model of cultural dissemination. Our main contribution has been the development of this model allowing us to isolate separate effects of these social forces, and the insights it has offered into the system dynamics arising from simple rules of social influence.

In Section 1.2 we hypothesised that this embedding of Social Impact Theory in a cultural model would *provide additional insights for synthetic problem scenarios while also aligning with data from real-world observations*. We have contributed a number of insights, perhaps the most striking of which has been the emergence of cultural overlap between agents; the *number* social force of Social Impact Theory allowing agents to adopt different traits from different neighbours without becoming completely identical to them. This stable cultural pluralism is enabled by social reinforcement, and Social Impact Theory combined with cultural dissemination offers a candidate explanation for the continued existence of overlapping cultures. In addition to uncovering these dynamics, appropriate metrics for measuring their prevalence are also contributed.

We have further examined the effects of different network structures on the dynamics of influence in *number* and over *distance*, allowing us to ascertain the role network characteristics such as node-degree and clustering play in behaviours of cultural models. A comparison is made with the theory of simple and complex contagions, and we determine local conditions which influence the mode of transmission in our model in similar - but not identical - ways to this theory.

Cultural mixing due to multiple competing influences has implications for the development of organisational culture. We undertook an analysis of trait spreading within hierarchical structures to examine the effects of the *hierarchical status* of superiors versus multiple subordinates in a manner which would not be possible in previous existing models. We have also contributed a method of generating generalised hierarchical structures incorporating 'team' sub-structures.

Agent-based models and social simulations are often theoretical and abstract in nature. However, in addition to general insights, we have conducted an empirical test of a selection of models on real world data, finding that models of multiple influence (such as ours) generally perform better than dyadic models. Notably, the polarisation many other models depict did not materialise in the real data.

8.1 Research questions

Here, we revisit questions posed in Chapter 1:

RQ1: Will trait copying based on *number* of influences produce different macroscopic results to the dyadic copying of Axelrod's model?

Where an agent takes into consideration the strength and immediacy of all of its influencers simultaneously, it can adopt different traits from each. This appears globally as a mixing and overlapping of different cultures. The complete polarisation of cultures - which was a defining characteristic of Axelrod's model - is no longer present. This establishes our model as one of cultural overlap (Chapter 3).

RQ2: Which local dynamics may offer explanations for global behaviours in our model that differ from those in Axelrod's?

The dyadic and stochastic trait-copying in Axelrod's model ensures that polarisation is inevitable. However, where the *number* factor of Social Impact Theory is incorporated into cultural contagion, an individual may receive social reinforcement from multiple sources, preventing assimilation (sections 3.3.1, 4.2).

RQ3: Can any differences between our model and Axelrod's be equated to the differences between complex and simple contagions?

The spread of traits in Axelrod's model aligns well with simple contagions, but in our model traits can spread as if by simple *or* complex contagion, depending on local context (Chapter 4). This supports the notion that simple and complex contagion may co-exist, and be applicable as consequence of the situation.

RQ4: What effect will network characteristics such as clustering have on the diffusion of traits using these models?

Increasing node-degree in Axelrod's model has a more marked effect on convergence than in ours. Axelrod's model seems more sensitive to average path length, whereas embedding Social Impact Theory leads to greater sensitivity to local clustering (Section 4.3.2).

RQ5: Where some agents are given increased levels of influence due to their *status*, will the shape of the hierarchy affect their ability to influence their subordinates?

The ways hierarchy affect trait spread are nuanced, and we found that upward influence from subordinates, even where individually weak, can have a notable effect on both cultural convergence and the ability of the higher nodes to disseminate their traits.

Horizontal links within 'teams' can serve to either undermine a superior's influence or to enhance it when the superior's strength is above a certain level.

RQ6: Will the addition of influence over structural *distance* (i.e. the *immediacy* of social contact) affect the core behaviours of cultural diffusion?

Increasing distance of influence tends to reduce the amount of cultural overlap in the population, but there are subtle differences in results from different methods of determining influence over distance. Varying clustering and path length within the network graph has a non-monotonic effect. An increase in small-world 'rewiring' initially promotes convergence at greater distance; however as the network becomes closer to a random graph this convergence can collapse into cultural fragmentation, due to the sheer number of equal competing influences.

RQ7: Given appropriate data, can computational culture models give an indication of how behaviour spreads through a real social network?

Most models give a reasonable approximation of the results from real data, with models of multiple influence performing best (Chapter 7).

8.2 Limitations and future work

We believe our work has provided new insights into some of the emergent patterns which result when applying Social Impact Theory to cultural modelling. However, there will always be limitations associated with modelling using simplifying assumptions, as no computational model can fully encompass every facet of human behaviour and social influence. There are several potential avenues for future investigation which could further refine our model and others in this area.

Many of our model's mechanisms are more deterministic than the existing models of social influence, although we retain a random order in the activation of agents. Whether this is more or less realistic is open to debate, but there is a tendency in existing stochastic models to *always* produce a state of polarisation. Where there is even the smallest chance of change, it will eventually occur; if the only limit to that change is a state of polarisation then such a result is inevitable. A future consideration could be a stochastic model with more open stop conditions, though this could lead to a state never reaching equilibrium.

An obvious extension would be the introduction of cultural drift, either through selection error or noise, as this is a feature of some existing models [84]. In previous models this had a tendency to encourage convergence to a monoculture; it may be interesting to see what effect our use of simultaneous *number* and *immediacy* forces from Social Impact Theory may have under such conditions.

The representation and categorisation of data is a broad problem, and one that has implications for the application of all social models. In the case of culture models, the characterisation of traits as being discrete and mutually exclusive has limitations; it assumes complete independence of traits with no overlap. In reality, some cultural or behavioural traits will be closer to others; for example, there are several distinct Romance languages but they often share much in common, and can justifiably be thought of as 'closer' to each other than to languages from another part of the world. Continuous values for opinions are also a simplification, as they represent opinion as a scalar value on a one-dimensional scale. A model that allows for continuous values, while still taking account of homophily and possible interacting contagions between features (i.e. the presence of a trait on one feature increases the chance of adoption of a similar trait on another) could be a possible extension. The possibility of *accumulating* traits rather than completely replacing existing ones is also a consideration, one which could combine models of cultural diffusion with those of cumulative culture [104].

When tested against real data, *distance* mechanisms appeared to perform less well than when only *strength* and *number* were applied (Chapter 7). The algorithms for determining influence over distance could be adjusted, either by scaling the effect up or down or using a different formulation. It could simply be that the effect is less prominent than we expected, and local contexts could make it more or less likely that referred influence is possible.

Finally, real-world data suitable for applying cultural dissemination models to remain scarce. Any additional datasets of similar nature to that of the 'Glasgow Teenage Friends' dataset would be a considerable boon to modelling in this area, and in the future we hope more models can be tested this way. A qualitative social study testing the results generated by our model would be ideal; however, it is potentially difficult to isolate these dynamics in the real world - which is one of the reasons for, and benefits of, adopting a modelling approach.

Bibliography

[1] AGNEW, R., AND PETERSEN, D. M. Leisure and delinquency. *Social problems 36*, 4 (1989), 332–350.

[2] AMBLARD, F., AND DEFFUANT, G. The role of network topology on extremism propagation with the relative agreement opinion dynamics. *Physica A: Statistical Mechanics and its Applications 343* (2004), 725–738.

[3] AMERICAN PSYCHOLOGICAL ASSOCIATION. APA dictionary of psychology. `https://dictionary.apa.org/social-influence`. Accessed: 2022-04-18.

[4] ARAZY, O., NOV, O., AND ORTEGA, F. The [wikipedia] world is not flat: On the organizational structure of online production communities. In *22nd European Conference on Information Systems, ECIS 2014* (2014), Association for Information Systems.

[5] ARENDT, H. *The origins of totalitarianism*, vol. 244. Houghton Mifflin Harcourt, 1973.

[6] AXELROD, R. The dissemination of culture: A model with local convergence and global polarization. *Journal of conflict resolution 41*, 2 (1997), 203–226.

[7] AXELROD, R., AND BENNETT, D. S. A landscape theory of aggregation. *British journal of political science 23*, 2 (1993), 211–233.

[8] AXELROD, R., DAYMUDE, J. J., AND FORREST, S. Preventing extreme polarization of political attitudes. *Proceedings of the National Academy of Sciences 118*, 50 (2021).

[9] AXTELL, R. L., EPSTEIN, J. M., DEAN, J. S., GUMERMAN, G. J., SWEDLUND, A. C., HARBURGER, J., CHAKRAVARTY, S., HAMMOND, R., PARKER, J., AND PARKER, M. Population growth and collapse in a multiagent model of the Kayenta Anasazi in Long House Valley. *Proceedings of the National Academy of Sciences 99*, suppl 3 (2002), 7275–7279.

[10] BAIL, C. A., ARGYLE, L. P., BROWN, T. W., BUMPUS, J. P., CHEN, H., HUNZAKER, M. F., LEE, J., MANN, M., MERHOUT, F., AND VOLFOVSKY, A. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences 115*, 37 (2018), 9216–9221.

[11] BAIL, C. A., GUAY, B., MALONEY, E., COMBS, A., HILLYGUS, D. S., MERHOUT, F., FREELON, D., AND VOLFOVSKY, A. Assessing the Russian Internet Research Agency's impact on the political attitudes and behaviors of American Twitter users in late 2017. *Proceedings of the National Academy of Sciences 117*, 1 (2020), 243–250.

[12] BAK, P., PACZUSKI, M., AND SHUBIK, M. Price variations in a stock market with many agents. *Physica A: Statistical Mechanics and its Applications 246*, 3-4 (1997), 430–453.

[13] BALDWIN, J. R., FAULKNER, S. L., HECHT, M. L., AND LINDSLEY, S. L. *Redefining culture: Perspectives across the disciplines*. Routledge, 2006.

[14] BALL, P. *Critical mass: How one thing leads to another*. Macmillan, 2004.

[15] BARABÁSI, A.-L., AND ALBERT, R. Emergence of scaling in random networks. *Science 286*, 5439 (1999), 509–512.

[16] BATTISTON, F., NICOSIA, V., LATORA, V., AND MIGUEL, M. S. Layered social influence promotes multiculturality in the Axelrod model. *Scientific Reports 7*, 1 (2017), 1–9.

[17] BAVELAS, A. Communication patterns in task-oriented groups. *The journal of the Acoustical Society of America 22*, 6 (1950), 725–730.

[18] BELLAMY, R., COLOMBO, G., ESHGHI, S., DE MEL, G., GIAMMANCO, C., MORRIS, R., RAND, D. G., TURNER, L. D., WHITAKER, R. M., AND WILLIAMS, G.-R. A computational framework for modelling inter-group behaviour using psychological theory. In *Next-Generation Analyst VI* (2018), vol. 10653, International Society for Optics and Photonics, p. 106530G.

[19] BERNSTEIN, E., BUNCH, J., CANNER, N., AND LEE, M. Beyond the holacracy hype. *Harvard Business Review 94*, 7 (2016), 8.

[20] BESSI, A., PETRONI, F., DEL VICARIO, M., ZOLLO, F., ANAGNOSTOPOULOS, A., SCALA, A., CALDARELLI, G., AND QUATTROCIOCCHI, W. Viral misinformation: The role of homophily and polarization. In *Proceedings of the 24th international conference on World Wide Web* (2015), pp. 355–356.

[21] BIANCHI, F., AND SQUAZZONI, F. Agent-based models in sociology. *Wiley Interdisciplinary Reviews: Computational Statistics 7*, 4 (2015), 284–306.

[22] BLACK, P. E. Dictionary of Algorithms and Data Structures. `https://xlinux.nist.gov/dads/HTML/manhattanDistance.html`. Accessed: 2022-10-30.

[23] BOUQUET, C., AND BIRKINSHAW, J. Managing power in the multinational corporation: How low-power actors gain influence. *Journal of Management 34*, 3 (2008), 477–508.

[24] BOUTYLINE, A., AND WILLER, R. The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political psychology 38*, 3 (2017), 551–569.

[25] Bowen, A. M., and Bourgeois, M. J. Attitudes toward lesbian, gay, and bisexual college students: The contribution of pluralistic ignorance, dynamic social impact, and contact theories. *Journal of American college health 50*, 2 (2001), 91–96.

[26] Box, G. E. Robustness in the strategy of scientific model building. In *Robustness in statistics*. Elsevier, 1979, pp. 201–236.

[27] Braines, D., Felmlee, D., Towsley, D., Tu, K., Whitaker, R. M., and Turner, L. D. The role of motifs in understanding behavior in social and engineered networks. In *Next-Generation Analyst VI* (2018), vol. 10653, International Society for Optics and Photonics, p. 106530W.

[28] Brown, A. R. Notes on the social organization of Australian tribes. *The Journal of the Royal Anthropological Institute of Great Britain and Ireland 48* (1918), 222–253.

[29] Burgelman, R. A. Strategy as vector and the inertia of coevolutionary lock-in. *Administrative Science Quarterly 47*, 2 (2002), 325–357.

[30] Burt, R. S. Structural holes. In *structural holes*. Harvard university press, 1992.

[31] Bush, H., West, P., and Michell, L. The role of friendship groups in the uptake and maintenance of smoking amongst pre-adolescent and adolescent children: Distribution of frequencies. *Glasgow, Scotland: Medical Research Council* (1997).

[32] Carley, K. M. Computational and mathematical organization theory: Perspective and directions. *Computational & mathematical organization theory 1*, 1 (1995), 39–56.

[33] Carzo Jr, R., and Yanouzas, J. N. Effects of flat and tall organization structure. *Administrative Science Quarterly* (1969), 178–191.

[34] CASTELLANO, C., FORTUNATO, S., AND LORETO, V. Statistical physics of social dynamics. *Reviews of modern physics 81*, 2 (2009), 591.

[35] CASTELLANO, C., MARSILI, M., AND VESPIGNANI, A. Nonequilibrium phase transition in a model for social influence. *Physical Review Letters 85*, 16 (2000), 3536.

[36] CENTOLA, D. *How behavior spreads*. Princeton University Press, 2018.

[37] CENTOLA, D., GONZALEZ-AVELLA, J. C., EGUILUZ, V. M., AND SAN MIGUEL, M. Homophily, cultural drift, and the co-evolution of cultural groups. *Journal of Conflict Resolution 51*, 6 (2007), 905–929.

[38] CENTOLA, D., AND MACY, M. Complex contagions and the weakness of long ties. *American journal of Sociology 113*, 3 (2007), 702–734.

[39] CIOFFI-REVILLA, C. Introduction to computational social science. *London and Heidelberg: Springer* (2014).

[40] COLEMAN, J. S. Social capital in the creation of human capital. *American journal of sociology 94* (1988), S95–S120.

[41] COLEMAN, J. S., KATZ, E., AND MENZEL, H. *Medical innovation: A diffusion study*. Indianapolis: Bobbs-Merrill Company, 1966.

[42] CONOVER, M., RATKIEWICZ, J., FRANCISCO, M., GONÇALVES, B., MENCZER, F., AND FLAMMINI, A. Political polarization on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media* (2011), vol. 5, pp. 89–96.

[43] DALEY, D. J., AND KENDALL, D. G. Epidemics and rumours. *Nature 204*, 4963 (1964), 1118–1118.

[44] DALL'ASTA, L., CASTELLANO, C., AND MARSILI, M. Statistical physics of the Schelling model of segregation. *Journal of Statistical Mechanics: Theory and Experiment 2008*, 07 (2008), L07002.

[45] DARAGANOVA, G., AND PATTISON, P. Social networks and space. In *International Workshop on Social Space and Geographic Space SGS'07* (2007), p. 3.

[46] DE SOLA POOL, I., AND KOCHEN, M. Contacts and influence. *Social networks 1*, 1 (1978), 5–51.

[47] DEFFUANT, G. Comparing extremism propagation patterns in continuous opinion models. *Journal of Artificial Societies and Social Simulation 9*, 3 (2006).

[48] DEFFUANT, G., AMBLARD, F., WEISBUCH, G., AND FAURE, T. How can extremism prevail? A study based on the relative agreement interaction model. *Journal of artificial societies and social simulation 5*, 4 (2002).

[49] DEFFUANT, G., NEAU, D., AMBLARD, F., AND WEISBUCH, G. Mixing beliefs among interacting agents. *Advances in Complex Systems 3*, 01n04 (2000), 87–98.

[50] DEL VICARIO, M., VIVALDO, G., BESSI, A., ZOLLO, F., SCALA, A., CALDARELLI, G., AND QUATTROCIOCCHI, W. Echo chambers: Emotional contagion and group polarization on facebook. *Scientific Reports 6*, 1 (2016), 1–12.

[51] DEMIRIS, N., KYPRAIOS, T., AND SMITH, L. V. On the epidemic of financial crises. *Journal of the Royal Statistical Society. Series A (Statistics in Society)* (2014), 697–723.

[52] DENISON, D. R. *Corporate culture and organizational effectiveness.* John Wiley & Sons, 1990.

[53] DIEFENBACH, T., AND SILLINCE, J. A. Formal and informal hierarchy in different types of organization. *Organization studies 32*, 11 (2011), 1515–1537.

[54] DORIGO, M., MANIEZZO, V., AND COLORNI, A. Ant system: optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics) 26*, 1 (1996), 29–41.

[55] DUNBAR, R. I. Neocortex size as a constraint on group size in primates. *Journal of human evolution 22*, 6 (1992), 469–493.

[56] ENGEL, U., QUAN-HAASE, A., LIU, S., AND LYBERG, L. *Handbook of Computational Social Science, Volume 1: Theory, Case Studies and Ethics*. European Association of Methodology Series. Taylor & Francis, 2021.

[57] EPSTEIN, J. M. Agent-based computational models and generative social science. *Complexity 4*, 5 (1999), 41–60.

[58] FAUST, K. A puzzle concerning triads in social networks: Graph constraints and the triad census. *Social Networks 32*, 3 (2010), 221–233.

[59] FLACHE, A., AND MACY, M. W. Local convergence and global diversity: From interpersonal to social influence. *Journal of Conflict Resolution 55*, 6 (2011), 970–995.

[60] FONTANARI, J. F. Social interaction as a heuristic for combinatorial optimization problems. *Physical Review E 82*, 5 (2010), 056118.

[61] FOSS, N. J., AND DOBRAJSKA, M. Valve's Way: Wayward, Visionary, or Voguish? *Journal of Organization Design 4*, 2 (2015), 12–15.

[62] FOTHERINGHAM, A. S. Spatial structure and distance-decay parameters. *Annals of the Association of American Geographers 71*, 3 (1981), 425–436.

[63] FRENCH, J. R., RAVEN, B., AND CARTWRIGHT, D. The bases of social power. *Classics of organization theory 7* (1959), 311–320.

[64] GALAM, S. Minority opinion spreading in random geometry. *The European Physical Journal B-Condensed Matter and Complex Systems 25*, 4 (2002), 403–406.

[65] GARBER, P. M. Tulipmania. *Journal of Political Economy 97*, 3 (1989), 535–560.

[66] GOFFMAN, W., AND NEWILL, V. Generalization of epidemic theory. *Nature 204*, 4955 (1964), 225–228.

[67] GONZÁLEZ-AVELLA, J. C., COSENZA, M. G., AND TUCCI, K. Nonequilibrium transition induced by mass media in a model for social influence. *Physical Review E 72*, 6 (2005), 065102.

[68] GRACIA-LÁZARO, C., LAFUERZA, L., FLORÍA, L., AND MORENO, Y. Residential segregation and cultural dissemination: An Axelrod-Schelling model. *Physical Review E 80*, 4 (2009), 046123.

[69] GRANOVETTER, M. S. The strength of weak ties. *American journal of sociology 78*, 6 (1973), 1360–1380.

[70] HANKINSON, P. An empirical study which compares the organisational structures of companies managing the World's top 100 brands with those managing outsider brands. *Journal of Product & Brand Management* (1999).

[71] HANSEN, M. T. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly 44*, 1 (1999), 82–111.

[72] HATCH, M. J., AND SCHULTZ, M. Relations between organizational culture, identity and image. *European Journal of marketing* (1997).

[73] HÉBERT-DUFRESNE, L., SCARPINO, S. V., AND YOUNG, J.-G. Macroscopic patterns of interacting contagions are indistinguishable from social reinforcement. *Nature Physics 16*, 4 (2020), 426–431.

[74] HEGSELMANN, R., KRAUSE, U., ET AL. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of artificial societies and social simulation 5*, 3 (2002).

[75] HILL, R. A., AND DUNBAR, R. I. Social network size in humans. *Human nature 14*, 1 (2003), 53–72.

[76] HOLLEY, R. A., AND LIGGETT, T. M. Ergodic theorems for weakly interacting infinite systems and the voter model. *The annals of probability* (1975), 643–663.

[77] ISARD, W. Location theory and trade theory: short-run analysis. *The Quarterly Journal of Economics* (1954), 305–320.

[78] JACKSON, J. M., AND LATANÉ, B. All alone in front of all those people: Stage fright as a function of number and type of co-performers and audience. *Journal of personality and social psychology 40*, 1 (1981), 73.

[79] JIN, F., DOUGHERTY, E., SARAF, P., CAO, Y., AND RAMAKRISHNAN, N. Epidemiological modeling of news and rumors on twitter. In *Proceedings of the 7th workshop on social network mining and analysis* (2013), pp. 1–9.

[80] KARAU, S. J., AND WILLIAMS, K. D. Social loafing: A meta-analytic review and theoretical integration. *Journal of personality and social psychology 65*, 4 (1993), 681.

[81] KERMACK, W. O., AND MCKENDRICK, A. G. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character 115*, 772 (1927), 700–721.

[82] KIPNIS, D., SCHMIDT, S. M., AND WILKINSON, I. Intraorganizational influence tactics: Explorations in getting one's way. *Journal of applied psychology 65*, 4 (1980), 440.

[83] Kirman, A. P., and Vriend, N. J. Evolving market structure: An ace model of price dispersion and loyalty. *Journal of Economic Dynamics and Control 25*, 3-4 (2001), 459–502.

[84] Klemm, K., Eguíluz, V. M., Toral, R., and San Miguel, M. Global culture: A noise-induced transition in finite systems. *Physical Review E 67*, 4 (2003), 045101.

[85] Klemm, K., Eguíluz, V. M., Toral, R., and San Miguel, M. Nonequilibrium transitions in complex networks: A model of social interaction. *Physical review E 67*, 2 (2003), 026120.

[86] Krend, J. *Reconstruction of Oliver Benson's "Simple Diplomatic Game"*. Inter-university Consortium for Political and Social Research, 1984.

[87] Kroeber, A. L., and Kluckhohn, C. Culture: A critical review of concepts and definitions. *Papers. Peabody Museum of Archaeology & Ethnology, Harvard University* (1952).

[88] Latané, B. The psychology of social impact. *American psychologist 36*, 4 (1981), 343.

[89] Latané, B., Liu, J. H., Nowak, A., Bonevento, M., and Zheng, L. Distance matters: Physical space and social impact. *Personality and Social Psychology Bulletin 21*, 8 (1995), 795–805.

[90] Lazarsfeld, P. F., Merton, R. K., et al. Friendship as a social process: A substantive and methodological analysis. *Freedom and control in modern society 18*, 1 (1954), 18–66.

[91] LeBon, G., and Nye, R. A. *The Crowd*. Routledge, 2017. (Originally published 1895).

[92] Levine, D. N., Carter, E. B., and Gorman, E. M. Simmel's influence on American sociology. i. *American Journal of Sociology 81*, 4 (1976), 813–845.

[93] LEWIS, K., KAUFMAN, J., GONZALEZ, M., WIMMER, A., AND CHRISTAKIS, N. Tastes, ties, and time: A new social network dataset using facebook.com. *Social networks 30*, 4 (2008), 330–342.

[94] LIM, B.-L., CHOI, M., AND PARK, M.-C. The late take-off phenomenon in the diffusion of telecommunication services: network effect and the critical mass. *Information economics and policy 15*, 4 (2003), 537–557.

[95] LUKE, D. A., AND HARRIS, J. K. Network analysis in public health: history, methods, and applications. *Annu. Rev. Public Health 28* (2007), 69–93.

[96] MACKAY, C. *Extraordinary popular delusions and the madness of crowds*. Simon and Schuster, 2012. (Original second edition published 1852).

[97] MAGEE, J., AND GALINSKY, A. Social hierarchy: The self-reinforcing nature of power and status. *Academy of Management annals 2*, 1 (2008), 351–398.

[98] MANZ, C. C., SHIPPER, F., AND STEWART, G. L. Everyone a team leader:: Shared influence at WL Gore & Associates. *Organizational dynamics 38*, 3 (2009), 239–244.

[99] MARCEAU, V., NOËL, P.-A., HÉBERT-DUFRESNE, L., ALLARD, A., AND DUBÉ, L. J. Modeling the dynamical interaction between epidemics on overlay networks. *Physical Review E 84*, 2 (2011), 026105.

[100] MARWELL, G., AND OLIVER, P. *The critical mass in collective action*. Cambridge University Press, 1993.

[101] MCADAM, D. Recruitment to high-risk activism: The case of freedom summer. *American journal of sociology 92*, 1 (1986), 64–90.

[102] MCADAM, D., AND PAULSEN, R. Specifying the relationship between social ties and activism. *American journal of sociology 99*, 3 (1993), 640–667.

[103] McPherson, M., Smith-Lovin, L., and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual review of sociology 27*, 1 (2001), 415–444.

[104] Mesoudi, A., and Thornton, A. What is cumulative cultural evolution? *Proceedings of the Royal Society B 285*, 1880 (2018), 20180712.

[105] Michell, L., and West, P. Peer pressure to smoke: the meaning depends on the method. *Health education research 11*, 1 (1996), 39–49.

[106] Milgram, S. The small world problem. *Psychology today 2*, 1 (1967), 60–67.

[107] Miller, J. H., Page, S. E., and Page, S. Complex adaptive systems. In *Complex Adaptive Systems*. Princeton university press, 2009.

[108] Mitchell, P. *Contagious metaphor*. A&C Black, 2012.

[109] Mobilia, M. Does a single zealot affect an infinite group of voters? *Physical review letters 91*, 2 (2003), 028701.

[110] Mobilia, M., and Georgiev, I. T. Voting and catalytic processes with inhomogeneities. *Physical Review E 71*, 4 (2005), 046102.

[111] Mobilia, M., Petersen, A., and Redner, S. On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment 2007*, 08 (2007), P08029.

[112] Morris, R., Turner, L., Whitaker, R., and Giammanco, C. The impact of peer pressure: Extending Axelrod's model on cultural polarisation. In *2019 IEEE International Conference on Cognitive Computing (ICCC)* (2019), IEEE, pp. 114–121.

[113] Morris, R. L., Turner, L. D., Whitaker, R. M., and Giammanco, C. Breadth verses depth: The impact of tree structure on cultural influence. In *Social, Cultural, and Behavioral Modeling* (Cham, 2020), R. Thomson, H. Bisgin,

C. Dancy, A. Hyder, and M. Hussain, Eds., Springer International Publishing, pp. 86–95.

[114] MORRIS, R. L., TURNER, L. D., WHITAKER, R. M., AND GIAMMANCO, C. The effect of hierarchical status and horizontal peer influence on organisational culture. Manuscript in preparation, 2022.

[115] MOSCOVICI, S. *Social influence and social change*, vol. 10. Academic Press, 1976.

[116] MOWDAY, R. T. The exercise of upward influence in organizations. *Administrative Science Quarterly* (1978), 137–156.

[117] MUNYON, T. P., SUMMERS, J. K., BROUER, R. L., AND TREADWAY, D. C. The implications of coalition forms for work role innovation, resource reallocation, and performance. In *Research in Personnel and Human Resources Management*. Emerald Group Publishing Limited, 2014.

[118] NEKOVEE, M., AND PINTO, J. Modeling the impact of organization structure and whistle-blowers on intra-organizational corruption contagion. *Physica A: Statistical Mechanics and its Applications 522* (2019), 339–349.

[119] NETTLE, D. Using social impact theory to simulate language change. *Lingua 108*, 2-3 (1999), 95–117.

[120] Networkx algorithms average clustering. `https://networkx.org/documentation/networkx-2.3/reference/algorithms/generated/networkx.algorithms.cluster.average_clustering.html#networkx.algorithms.cluster.average_clustering`. Accessed: 2021-10-26.

[121] Networkx generators connected_watts_strogatz_graph. `https://networkx.org/documentation/networkx-2.3/reference/generated/networkx.generators.random_graphs.connected_watts_strogatz_graph.html#`

`networkx.generators.random_graphs.connected_watts_strogatz_graph`. Accessed: 2021-10-26.

[122] NOWAK, A., SZAMREJ, J., AND LATANÉ, B. From private attitude to public opinion: A dynamic theory of social impact. *Psychological review 97*, 3 (1990), 362.

[123] OEDZES, J. J., VAN DER VEGT, G. S., RINK, F. A., AND WALTER, F. On the origins of informal hierarchy: The interactive role of formal leadership and task complexity. *Journal of Organizational Behavior 40*, 3 (2019), 311–324.

[124] OLSSON, G. Distance and human interaction a migration study. *Geografiska Annaler: Series B, Human Geography 47*, 1 (1965), 3–43.

[125] ONSAGER, L. Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Physical Review 65*, 3-4 (1944), 117.

[126] PARKER, M. *Organizational culture and identity: Unity and division at work*. Sage, 2000.

[127] PARKHURST, J. O., CHILONGOZI, D., AND HUTCHINSON, E. Doubt, defiance, and identity: understanding resistance to male circumcision for HIV prevention in Malawi. *Social Science & Medicine 135* (2015), 15–22.

[128] PASTOR-SATORRAS, R., CASTELLANO, C., VAN MIEGHEM, P., AND VESPIGNANI, A. Epidemic processes in complex networks. *Reviews of modern physics 87*, 3 (2015), 925.

[129] PEARSON, M., AND MICHELL, L. Smoke rings: social network analysis of friendship groups, smoking and drug-taking. *Drugs: education, prevention and policy 7*, 1 (2000), 21–37.

[130] PEARSON, M., AND WEST, P. Drifting smoke rings: Social network analysis and Markov processes in a longitudinal study of friendship groups and risk-taking. *Connections 25*, 2 (2003), 59–76.

[131] PRISLIN, R., AND CRANO, W. D. A history of social influence research. *Handbook of the History of Social Psychology* (2012).

[132] PUTNAM, R. D. Political attitudes and the local community. *American Political Science Review 60*, 3 (1966), 640–654.

[133] RAPOPORT, A. Contribution to the theory of random and biased nets. *The bulletin of mathematical biophysics 19*, 4 (1957), 257–277.

[134] REILLY, W. J. *The law of retail gravitation*. WJ Reilly, 1931.

[135] RENTFROW, P. J., AND GOSLING, S. D. The do re mi's of everyday life: the structure and personality correlates of music preferences. *Journal of personality and social psychology 84*, 6 (2003), 1236.

[136] SCHEIN, E. H. *Organizational culture.*, vol. 45. American Psychological Association, 1990.

[137] SCHELLING, T. C. Dynamic models of segregation. *Journal of mathematical sociology 1*, 2 (1971), 143–186.

[138] SCHELLING, T. C. *Micromotives and macrobehavior*. WW Norton & Company, 2006. (Originally published 1978).

[139] SCHULT, D. A. Exploring network structure, dynamics, and function using networkx. In *In Proceedings of the 7th Python in Science Conference (SciPy* (2008), Citeseer.

[140] SHIBANAI, Y., YASUNO, S., AND ISHIGURO, I. Effects of global information feedback on diversity: extensions to Axelrod's adaptive culture model. *Journal of Conflict Resolution 45*, 1 (2001), 80–96.

[141] SIGHELE, S. The criminal crowd and other writings on mass society. In *The Criminal Crowd and Other Writings on Mass Society*. University of Toronto Press, 2019. (Originally published 1892).

[142] SIMMEL, G. *The sociology of Georg Simmel*, vol. 92892. Simon and Schuster, 1950.

[143] SINGH, P., SREENIVASAN, S., SZYMANSKI, B. K., AND KORNISS, G. Accelerating consensus on coevolving networks: The effect of committed individuals. *Physical Review E 85*, 4 (2012), 046104.

[144] SNIJDERS, T. A. Excerpt of 50 girls from 'Teenage friends and lifestyle study' data. `https://www.stats.ox.ac.uk/~snijders/siena/s50_data.htm`. Accessed: 2022-02-21.

[145] SNIJDERS, T. A. 'Teenage friends and lifestyle study' data. `https://www.stats.ox.ac.uk/~snijders/siena/Glasgow_data.htm`. Accessed: 2022-02-21.

[146] SOLOMONOFF, R., AND RAPOPORT, A. Connectivity of random nets. *The bulletin of mathematical biophysics 13*, 2 (1951), 107–117.

[147] SPICER, A. No bosses, no managers: the truth behind the 'flat hierarchy' facade. *The Guardian 30* (2018).

[148] STEINERT-THRELKELD, Z. C. Spontaneous collective action: Peripheral mobilization during the Arab spring. *American Political Science Review 111*, 2 (2017), 379–403.

[149] STEWART, J. Q. An inverse distance variation for certain social influences. *Science 93*, 2404 (1941), 89–90.

[150] STOCKER, R., CORNFORTH, D., AND BOSSOMAIER, T. R. Network structures and agreement in social network simulations. *Journal of Artificial societies and social simulation 5*, 4 (2002).

[151] TELESFORD, Q. K., JOYCE, K. E., HAYASAKA, S., BURDETTE, J. H., AND LAURIENTI, P. J. The ubiquity of small-world networks. *Brain connectivity 1*, 5 (2011), 367–375.

[152] TRAVERS, J., AND MILGRAM, S. An experimental study of the small world problem. *Sociometry 32*, 4 (1969), 425–443.

[153] URBACH, T., AND FAY, D. When proactivity produces a power struggle: How supervisors' power motivation affects their support for employees' promotive voice. *European Journal of Work and Organizational Psychology 27*, 2 (2018), 280–295.

[154] VALVE CORPORATION. Handbook for new employees. `https://cdn.akamai.steamstatic.com/apps/valve/Valve_NewEmployeeHandbook.pdf`, 2012. [Online; accessed 6-July-2021].

[155] VAN DER STRATEN, A., MONTGOMERY, E. T., MUSARA, P., ETIMA, J., NAIDOO, S., LABORDE, N., HARTMANN, M., LEVY, L., BENNIE, T., CHENG, H., ET AL. Disclosure of pharmacokinetic drug results to understand nonadherence: results from a qualitative study. *AIDS (London, England) 29*, 16 (2015), 2161.

[156] VICARIO, M. D., QUATTROCIOCCHI, W., SCALA, A., AND ZOLLO, F. Polarization and fake news: Early warning of potential misinformation targets. *ACM Transactions on the Web (TWEB) 13*, 2 (2019), 1–22.

[157] WANG, L., AND WOOD, B. C. An epidemiological approach to model the viral propagation of memes. *Applied Mathematical Modelling 35*, 11 (2011), 5442–5447.

[158] WARK, C., AND GALLIHER, J. F. Emory Bogardus and the origins of the social distance scale. *The American Sociologist 38*, 4 (2007), 383–395.

[159] WATTS, D. J. *Small worlds*. Princeton university press, 2018.

[160] WATTS, D. J., AND STROGATZ, S. H. Collective dynamics of 'small-world' networks. *Nature 393*, 6684 (1998), 440–442.

[161] WORTHY, J. C. Organizational structure and employee morale. *American Sociological Review 15*, 2 (1950), 169–179.

[162] XIE, J., EMENHEISER, J., KIRBY, M., SREENIVASAN, S., SZYMANSKI, B. K., AND KORNISS, G. Evolution of opinions on social networks in the presence of competing committed groups. *PLoS One 7*, 3 (2012), e33215.

[163] YIN, Z., KATIMS, D. S., AND ZAPATA, J. T. Participation in leisure activities and involvement in delinquency by mexican american adolescents. *Hispanic Journal of Behavioral Sciences 21*, 2 (1999), 170–185.

[164] YUKL, G., AND FALBE, C. M. Influence tactics and objectives in upward, downward, and lateral influence attempts. *Journal of applied psychology 75*, 2 (1990), 132.

[165] ZIMMER, M. "But the data is already public": on the ethics of research in facebook. *Ethics and information technology 12*, 4 (2010), 313–325.

[166] ZIPF, G. *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*. Addison-Wesley Press, 1949.

Appendices

Appendix A

Additional results of influence over distance

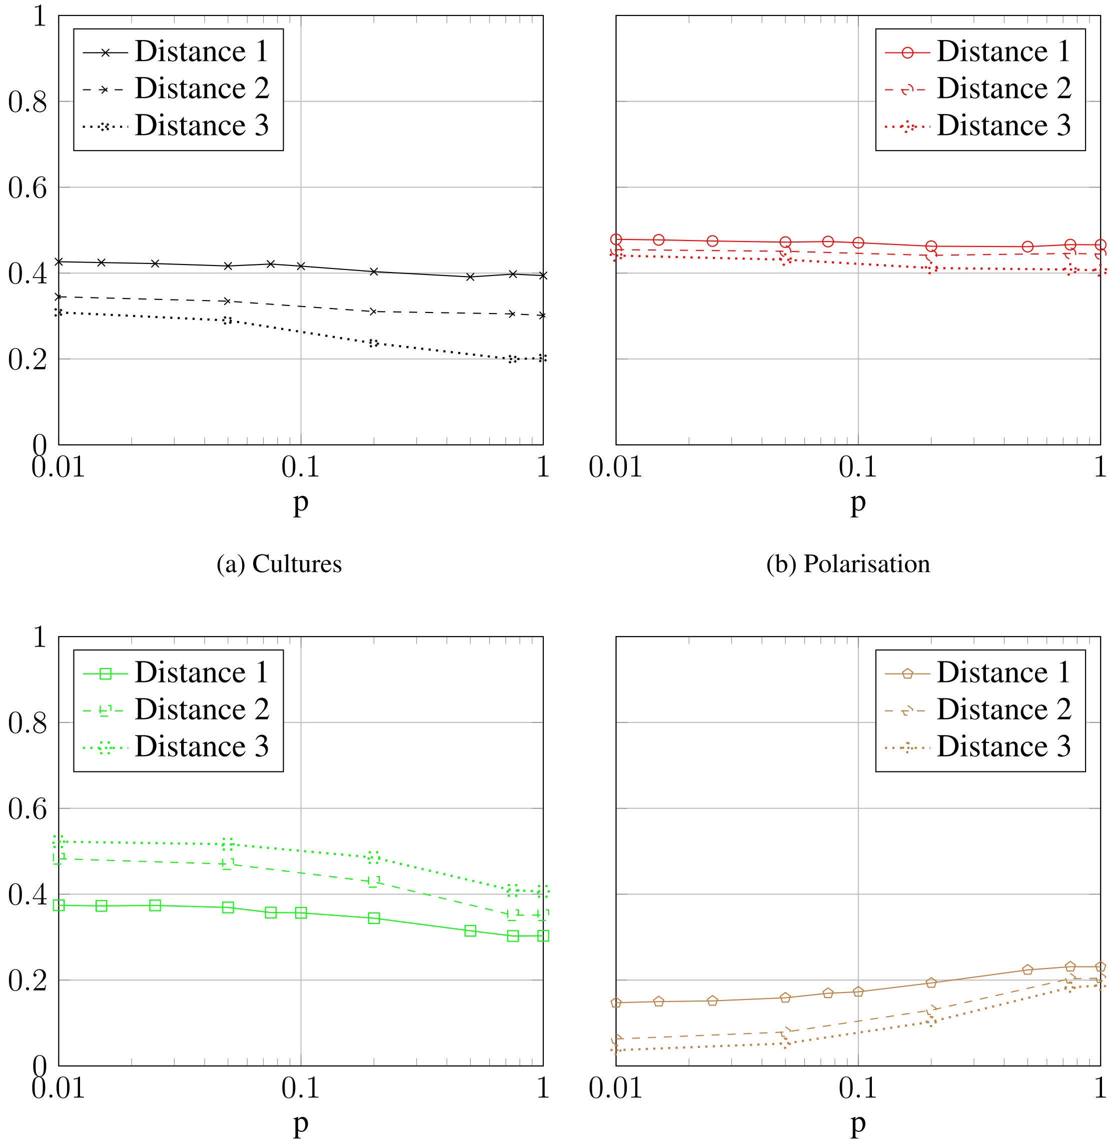

(a) Cultures (b) Polarisation

(c) Homogenisation (d) Mixing

Figure A.1: Results of path-based distance in Watts-Strogatz small-world networks ($degree$ = 4): At low degree the results patterns are largely consistent with those where no distance-mechanic is applied.

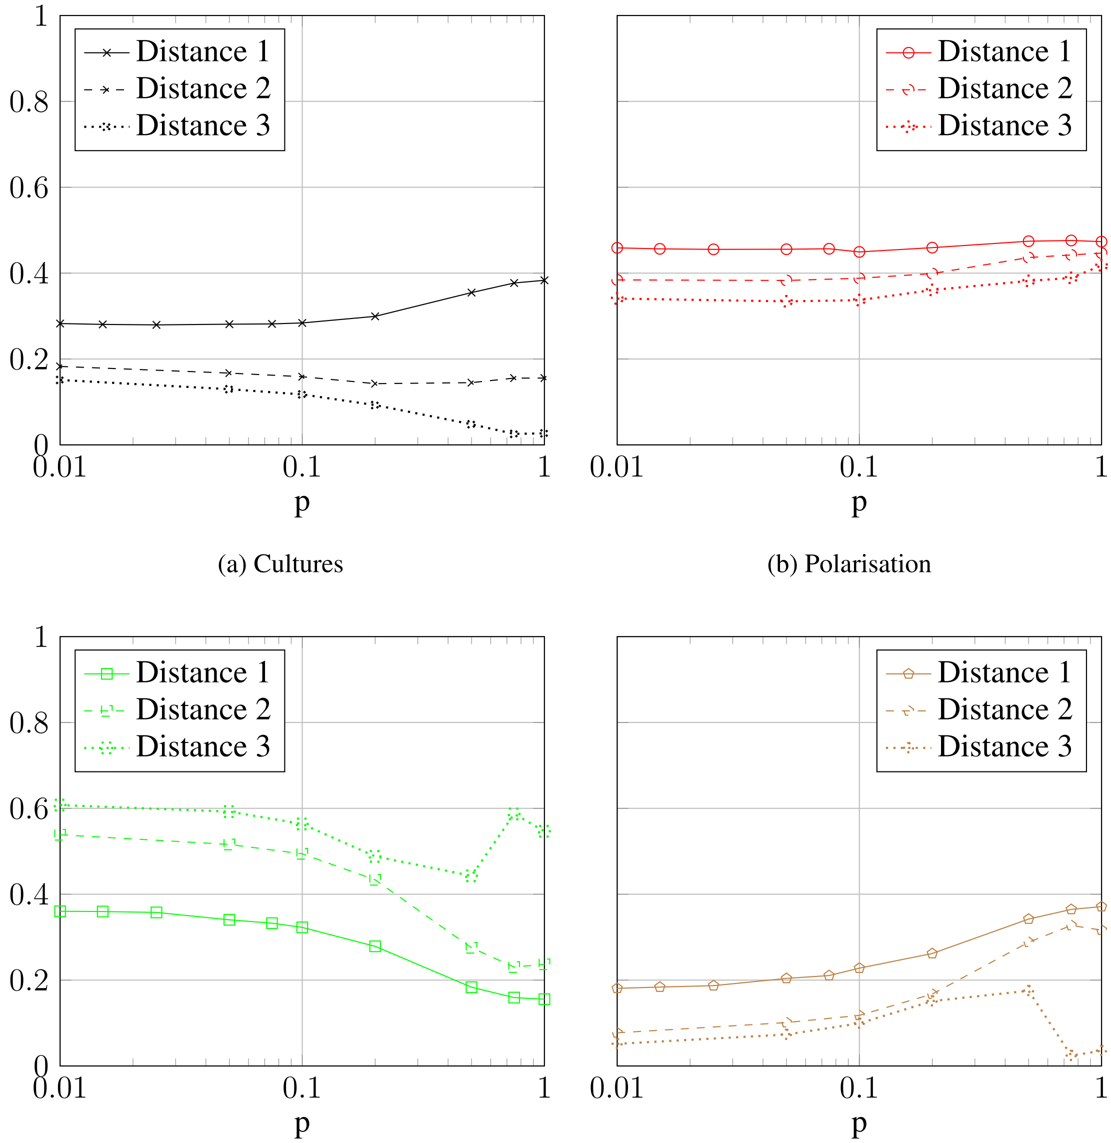

(a) Cultures

(b) Polarisation

(c) Homogenisation

(d) Mixing

Figure A.2: Results of path-based distance in Watts-Strogatz small-world networks ($degree$ = 8**).**

Appendix B

Glasgow Dataset Fields

Attribute	Values	Mutable	Longitudinal	Note
Friendship	Adjacency matrix of pupil ids and relationships; 0 = not a friend, 1 = "best friend", 2 = "just a friend"	Yes	Yes	
Romantic	1 (No), 2 (Yes)	Yes	Yes	In romantic relationship.
selection129	TRUE, FALSE	No	No	Whether pupil is in set of pupils who responded to survey at all three timesteps.
Tobacco	1 (non), 2 (occasional), 3 (regular)	Yes	Yes	
Alcohol	1 (non), 2 (once or twice a year), 3 (once a month), 4 (once a week), 5 (more than once a week)	Yes	Yes	
Cannabis	1 (non), 2 (tried once), 3 (occasional), 4 (regular)	Yes	Yes	
Rock	0, (No), 1 (Yes)	Yes	Yes	
Chart	0, (No), 1 (Yes)	Yes	Yes	
Reggae	0, (No), 1 (Yes)	Yes	Yes	
Dance	0, (No), 1 (Yes)	Yes	Yes	

Attribute	Values	Mutable	Longitudinal	Note
Heavy Metal	0, (No), 1 (Yes)	Yes	Yes	
Techno	0, (No), 1 (Yes)	Yes	Yes	
Folk/traditional	0, (No), 1 (Yes)	Yes	Yes	
Rave	0, (No), 1 (Yes)	Yes	Yes	
Indie	0, (No), 1 (Yes)	Yes	Yes	
Jazz	0, (No), 1 (Yes)	Yes	Yes	
Classical	0, (No), 1 (Yes)	Yes	Yes	
60s/70s	0, (No), 1 (Yes)	Yes	Yes	
House	0, (No), 1 (Yes)	Yes	Yes	
Grunge	0, (No), 1 (Yes)	Yes	Yes	
Rap	0, (No), 1 (Yes)	Yes	Yes	
Hip-hop	0, (No), 1 (Yes)	Yes	Yes	
I listen to tapes or CDs	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I look around in the shops	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	

Attribute	Values	Mutable	Longitudinal	Note
I read comics, mags or books	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I go to sport matches	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I take part in sports	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I hang round in the streets	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I play computer games	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I spend time on my hobby (eg art, an instrument)	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I go to something like B.B., Guides or Scouts	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	

Attribute	Values	Mutable	Longitudinal	Note
I go to cinema	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I go to pop concerts, gigs	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I go to church, mosque or temple	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I look after a pet animal	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I go to dance clubs or raves	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
I do nothing much (am bored)	1 (most days), 2 (once a week), 3 (once a month), 4 (less often or never)	Yes	Yes	
Distance to school	Numeric	Yes	Yes	
Distance to other pupils	Matrix of distances to other pupils	Yes	Yes	

Attribute	Values	Mutable	Longitudinal	Note
Angle	Matrix of angles	Yes	Yes	Angle between the lines connecting two students' homes with the school, given as a matrix.
Money	Numeric	Yes	Yes	Pocket money per month
Smoking at home	1 (No), 2 (Yes)	No	No	
Smoking by at least one parent	2 (No), 2 (Yes)	No	No	
Smoking by at least one sibling	3 (No), 2 (Yes)	No	No	
Age	Numeric	No	No	
Sex	1 (male), 2 (female)	No	No	

Table B.1: Full list of fields from the ‘Teenage friends and lifestyle study’ dataset.