



Provided by the author(s) and University of Galway in accordance with publisher policies. Please cite the published version when available.

Title	The influence of random interactions and decision heuristics on norm evolution in social networks
Author(s)	Mungovan, Declan; Howley, Enda; Duggan, Jim
Publication Date	2011-05
Publication Information	Mungovan, D,Howley, E,Duggan, J (2011) 'The influence of random interactions and decision heuristics on norm evolution in social networks'. Computational And Mathematical Organization Theory, 17 :152-178.
Publisher	Springer
Link to publisher's version	<a href="http://dx.doi.org/10.1007/s10588-011-9085-7">http://dx.doi.org/10.1007/s10588-011-9085-7</a>
Item record	<a href="http://hdl.handle.net/10379/3934">http://hdl.handle.net/10379/3934</a>
DOI	<a href="http://dx.doi.org/DOI 10.1007/s10588-011-9085-7">http://dx.doi.org/DOI 10.1007/s10588-011-9085-7</a>

Downloaded 2024-04-18T04:44:58Z

Some rights reserved. For more information, please see the item record link above.



# The Influence of Random Interactions and Decision Heuristics on Norm Evolution in Social Networks

Received: date / Accepted: date

**Abstract** In this paper we explore the effect that random social interactions have on the emergence and evolution of social norms in a simulated population of agents. In our model agents observe the behaviour of others and update their norms based on these observations. An agent's norm is influenced by both their own fixed social network plus a second random network that is composed of a subset of the remaining population. Random interactions are based on a weighted selection algorithm that uses an individual's path distance on the network to determine their chance of meeting a stranger. This means that friends-of-friends are more likely to randomly interact with one another than agents with a higher degree of separation. We then contrast the cases where agents make highest utility based completely rational decisions about which norm to adopt versus using a Markov Decision process that associates a weight the best choice. Finally we examine the effect that these random interactions have on the evolution of a more complex social norm as it propagates throughout the population. We discover that increasing the frequency and weighting of random interactions results in higher levels of norm convergence and in a quicker time when agents have the choice between two competing alternatives. This can be attributed to more information passing through the population thereby allowing for quicker convergence. When the norm is allowed to evolve we observe both global consensus formation and group splintering depending on the cognitive agent model used.

## 1 Introduction

Social norms, or normative behaviours, are one mechanism that allows large groups of self interested humans to cooperation together and coordinate actions (Lopez y Lopez *et al.*, 2006) thereby providing a solution to the problem

---

Address(es) of author(s) should be given

of social order (Horne, 2007). They can be defined as a set of conventions or behavioural expectations that people in a population abide by. Essentially norms inform an agent, or individual, on how to behave. Ignoring social norms can lead to negative repercussions for individuals including being excluded from a group. Social norms present a balance between individual freedom on the one hand and the goals of the society on the other (Walker and Woolridge, 1995). Indeed the agency versus structure debate has been identified as an important questions in social science (Shilling, 1999). It aims to understand the intrinsic motivations of how humans make decisions. Is it through the structural institutions we encounter or our own intrinsic internal decision making mechanisms? Reconciling both ideas in the theory of structuration proposed by Giddens (1984). It states that human actions are performed within social structures including norms but these social norms themselves are transient and are prone to evolution. There are two types of social norm conventions: top-down and bottom up. Top-down norms represent laws that are enforced on the population (Goldfarb and Henrekson, 2003). Bottom up conventions, such as shaking hands when introducing oneself, represents emergent behaviour from within the group. In this scenario agents, acting in their own self interest, choose which action to take based upon their interactions with others in the population. This type of “conformist transmission” is a tendency to adopt the most popular behaviour in the group and aids in the convergence of social norms (Henrick and Boyd, 2001). It can be seen as a form of herding (Barrat, 2008) and is the type of social norm that we investigate in this paper. Norms share many characteristics in common with epidemic diseases which spread both horizontally and obliquely (Klein, 1999). Agents use locally available information to determine their selection of social norms. Social evolution in turn can be described as changes in the nongenetic information stored in societies (Ehrlich and Levin, 2005). In other words, over time the characteristics of a norm that describe its behaviour can change.

The model we present in this paper investigates the effects of random social contacts on norm convergence and evolution. In the real world individuals are unlikely to change their immediate social network of acquaintances very much but will have a number of once-off interactions with random members of the population. An individual will, generally speaking, have the same wife, boss, friend etc. from one day to the next. But it is the random meetings with the general public that we interested in. For instance we interact with complete strangers on buses, in shops, at parties etc. These interactions are outside ones social network. The interactions don't necessarily need to be verbal or significant. Norms have the capacity to influence peoples behaviour just through observation. We recognise, however, that because our daily activities are tied to our social network, we are more likely to randomly meet some people over others. To account for this we bias random interactions based on the social distance that separates agents in the network. We contrast two separate mechanisms of how agents make decisions. Firstly agents act in a perfectly rational manner based on the utility they perceive for each norm. Secondly we use a Markov Decision Process where an agent's choice is weighted toward making a

---

particular choice given their observations on the network. We then expand the model from a simple binary decision between two separate norms to incorporate a continuum that describes the norm as a series of  $n$  bits. This examines how the characteristics of the norm evolves over time given that agents are randomly interacting with one another. Our goal is to discover at what point random interactions will influence the emergence of a global convention and how these random interactions effect the evolution of the norm. Specifically, we we aim to:

1. Design an algorithm that selects a random individual based on their social distance in the network.
2. Test the effects of random interactions on both perfectly rational agents and agents that use a Markov Decision Process that select between two competing norms.
3. Test the effect of random interactions on an evolving norm that can change over time.

The formal model we describe, and the decision making algorithms we use, have been abstracted from real human behaviour into a multi-agent based simulation. Social life has many more complications but it can be useful to consider simpler settings (Horne, 2007). Agent Based Modelling (ABM)<sup>1</sup> often has the function of identifying new questions as highlighted by Epstein (2008). One advantage of our reduction in sophistication or veridicality is that it allows for an increased population of agents. As Carley (2009) points out, models are often developed for the purpose of telling a story or making a point and are not always meant for developing policy or guiding decisions. In this spirit our work is not accompanied by a real world statistical comparison. Nevertheless, our approach offers novel insight into the process of norm convergence and evolution. The rest of the paper is structured as follows; Section 2 presents an introduction to previous work in the area of norm convergence, evolution and social networks. Section 3 gives a description of the formal model used to define the agent based simulator and an explanation of how the simulator was designed and implemented. Section 4 presents the experimental results. Finally, in Section 5 we outline our conclusions and possible future work.

## 2 Related Research

We now explore the existing literature as it relates to our research. First we introduce the domain of theoretical normative systems. We then expand this to describe ABMs of simple norm convergence. This is followed by a look at cultural evolution and opinion dynamics in ABMs that have had an influence on our own model design. We follow this with an analysis of various social networks and their metrics. Networks have been shown to play an important role on dynamic social processes. We conclude the section with a brief overview

---

<sup>1</sup> In this paper ABM stands for both Agent Based Modelling and Agent Based Model depending on the context of its use.

of dynamic networks and some insights into what we mean by randomness in ABMs.

## 2.1 Normative behaviour and Social Dynamics

Recent studies by Loeffbroer et al. (2009) have demonstrated the importance of social norms in societies where the individualisation process is fairly advanced. They also observe the formation and emergence of new norms that replace older ones. Bikhchandani et al. (1992; 1998) describe a decision making model where individuals update their own belief system based on the observation of others. They note the formation of what they describe as an “information cascade” whereby a particular piece of information is adopted by everyone regardless of its intrinsic merit.

Agent-based Modelling has been used in recent years as a method of studying these social norms (Mukherjee et al., 2008; Lopez y Lopez et al., 2006; Conte et al., 1999; Villatoro et al., 2009). Conte et al. (1999) and Walker et al. (1995) describe a framework on integrating concepts of Multi Agent Systems with normative behaviour and how both disciplines interact. Using ABM Lopez-Pintado et al. (2008) describe social influence in the context of norms that make a binary decision<sup>2</sup>. They describe an influence response function that assigns a weighted number to the alternatives. Others such as Bikhchandani et al. (1998; 1992), Centola et al. (2005) and Watts et al. (2007) have also taken this approach of defining the adoption or diffusion of norms within a population as a choice between two competing alternatives. In these models an agent observes the choices of others and is influenced by their decisions.

But what if we want to describe the norm as something more complex than a choice between two competing alternatives? Norms themselves can change due to alterations made by the participants. Individuals can modify different aspects of the norm to suit their own needs (Ostrom, 2000). There are a number of models in the domain of opinion dynamics and the evolution of cultural features that are closely related to norm convergence and evolution that have a bearing on our research. Axelrod’s (1997) model on the dissemination of culture on a lattice network investigates how cultural traits can be influenced by interacting agents. The state of an agent is defined by its cultural traits which take the form of  $F$  components and  $q$  different values. A cultural component is a unique cultural characteristic such as music preference which can have any of  $q$  different values. We can see then that the total spectra of possible unique cultural states of an agent equals  $q^F$ . Agents will interact with one another if there is enough cultural overlap. Others such as Gonzalez-Avella et al. (2005), Kuperman (2006) and Centola et al. (2007) have extended this idea to incorporate other features. In these papers, as in the Axelrod (1997) model, agents interact with a *probability* based on their cultural overlap. The higher the overlap, the greater chance of influence.

---

<sup>2</sup> When we say binary decision we mean the choice between two competing alternatives.

---

Next we describe a model of consensus formation which also influence our model design of norm evolution. Deffuant et al.(2001), Kozma et al. (2008) and McKeown et al. (2006) describe a consensus formation model where agents can take on an opinion in the range of between 0 and 1. The opinion of interacting agents is then influenced by their neighbours on an Erdős and Rényi random graph. As in the models described above, (Axelrod, 1997; Gonzalez-Avella *et al.*, 2005; Kuperman, 2006), if the opinion between two agents is too far away then there is no influence. However the notion of a bounded confidence is described by a tolerance parameter,  $d$ , that can be tuned to lead to different states of polarisation. Once there is sufficient cultural overlap then the agents will *always interact*. In these cases a position halfway between both of their states is adopted.

In our models described in Sections 3.4 and 3.5 we use this concept of similarity influence but relate it to social norms rather than the dissemination of culture. In other words, if an agent observes a norm that is too distant from their own then they will not be influenced by it. Conversely, if an agent observes a social norm that is close enough to the one they currently employ they will alter their own  $n$ -bit norm to become more closely related. A lot of social science research supports the idea of homophily when describing human cultural traits (McPherson *et al.*, 2001; Ruef *et al.*, 2003). Individuals will change their views, norms etc. to become more like those they interact with. We have opened this literature review by discussing generally issues relating to normative behaviour. We then went on to describe some models using agent based approaches where agents are presented with a binary decision of two competing norms. We then follow this with an analysis of evolving cultural traits and dynamic opinion models. These are concepts we shall use later on when describing our own framework. In the next section we will look at social networks and how they effect the dissemination of norms and information.

## 2.2 Social Networks

A considerable amount of the literature has studied the effects of norm emergence in populations that are fully connected or interact in a random fashion (Shoham and Tennenholtz, 1995; Walker and Woolridge, 1995). The network that agents interact on, however, has been shown to play a significant role on the dynamics of diffusion (Kittock, 1995; Villatoro *et al.*, 2009; Savarimuthu *et al.*, 2007; Rahmandad and Sterman, 2008). Most of this work has dealt with static networks that are generated at initialisation time and do not change for the duration of the simulation. There have, however, been attempts to frame research within the bounds of dynamic networks (Savarimuthu *et al.*, 2007). Networks can be envisioned as a series of nodes<sup>3</sup>,  $N$ , that each have  $k$  links to other nodes. Depending on the network type,  $k \geq 0$ . In the next section we give an overview of different network types that agents can interact on. We

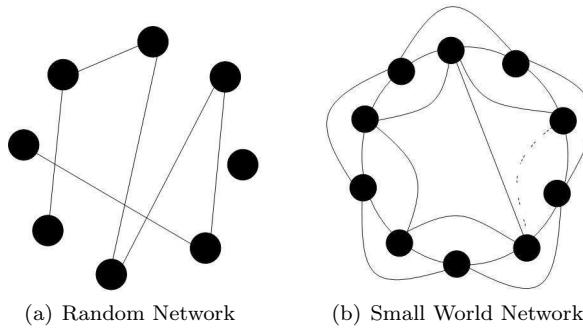
---

<sup>3</sup> Keeping with the nomenclature used in the literature we use the terms Node and Vertex interchangeable

are now going to look at two different types of networks that are important to our research. The first of these is one of the most widely studied graphs in network theory (Erdos and Reyni, 1961), the random network.

### 2.2.1 Random Network

The *Random Graph* of Erdős and Rényi first gained prominence in the 1950s. However, random graphs differ from real world structures in that they lack clustering or transitivity (Newman, 2002). A Random Network consists of  $N$  nodes randomly connected to one another. The probability of a connection between two nodes is fixed where:  $p_c = k/(N - 1)$  and  $k$  = average number of links per person. Figure 1(a) shows an example of a simple random network where  $k = 1.5$  and  $N = 8$  and  $p_c = 0.21$ . Notice that it is possible for some nodes on the network to be completely disconnected from the rest. Next we look at Small World Networks that more closely represent networks found in the real world.



**Fig. 1**

### 2.2.2 Small World Social Networks

The idea of *Small World (SW) Networks* first gained widespread popularity with Stanley Milgram's (1967) small-world study of large and sparse networks. Watts et al. (1998) later describe these networks as being formed by rewiring the edges of regular lattices with probability  $p_w$ . SW Networks are highly clustered, yet have length scaling properties equivalent to the expectations of randomly assembled graphs (Watts, 1999a). Notice in Fig. 1(b) that the link with the dashed line has been re-wired to another part of the network. This creates an instant shortcut to distant nodes. SW graphs span the gap between ordered lattices and random graphs. Note that when  $p_w = 1$ , then all links are randomly assigned and the network becomes a random network. Lee et al. (2006) investigated the effect that changing the value of  $p_w$  has on the emergence of a winner take all outcome in product adoption. They discovered

---

that as  $p_w$  is increased the chance of a winner take all outcome becomes more likely. This is because as the value of  $p_w$  gets closer to one the network starts to become more like a random network. This prevents localized cliques of products from existing. An analysis of a number of real world human networks (Watts, 1999b; Verspagen and Duysters, 2004; Davis *et al.*, 2003; Baum *et al.*, 2003) have shown that they form SW networks.

We have just looked at some general properties of both SW and Random Networks. Next we will discuss some of the metrics that are used to analyse networks that help reveal the underlying properties of their nature.

### *2.2.3 Clustering Coefficient*

The clustering coefficient of a vertex in a network is a measure of the proportion of neighbours of that vertex that are neighbours of one another (Watts, 1999a). The more sparse and random a network is the lower the clustering. For node  $i$  the value of the clustering coefficient,  $c_i$ , is calculated as described in equation 1 where  $\varepsilon_i$  is the total number of edges between the neighbours of node  $i$  and  $k_i$  is the total number of neighbours of node  $i$ .

$$c_i = \frac{2\varepsilon_i}{k_i(k_i - 1)} \quad (1)$$

Next we describe the path length which is another important network property used in analysing social networks.

### *2.2.4 Path Length*

The Maximum Path Length (MPL) is the maximum number of steps required to get to the furthest node, or nodes, on the network. Dijkstra's algorithm (Zhan and Noon, 1998) uses a breadth first search to traverse the network and discover the shortest path to each agent. The average shortest path length is the average number of steps along the shortest paths for all possible pairs of network nodes. It is a measure of the how quickly information etc. can diffuse through the network. Studies have shown that random networks and SW networks have a low average path length (Fronczak *et al.*, 2002; Watts, 1999a). So far we have discussed norms in agent based models along with network theory and some of the important metrics used to analyse them. The SW and Random networks of those studies have involved static configurations that don't change. We will now mention research that utilises dynamic networks where links/ nodes are changing.

### *2.2.5 Dynamic Networks*

Fenner *et al.* (2007) describe a stochastic model for a social network. Individuals may join the network, existing actors may become inactive and, at a later stage, reactivate themselves. The model captures the dynamic nature

of the network as it evolves over time. Actors attain new relations according to a preferential attachment rule that weights different agents according to their degree<sup>4</sup>. Savarimuthu et al. (2007) generated a network of agents whose topology changes dynamically. Agents initially randomly collide on a 2D grid and then proceed to form social networks. Kozma et al. (2008) investigate the effects of consensus formation using a dynamical network topology that allows agents to rewire links. We define what we mean by randomness when we speak of agents in a simulated computer environment.

### 2.3 Randomness in Agent Based Modeling

Agent based modelling has a number of advantages over classical game theory approaches. Firstly, ABMs are capable of implementing Monte Carlo<sup>5</sup> type stochastic iterations of a complex system. Izquierdo et al. (2009) highlight the fact that any computer model is in fact, due to its very nature, deterministic. However, we can use pseudo-random number generators to simulate random variables within the model and generate an artificial Monte Carlo generator. The pseudo-random number generator is an algorithm that takes in a random input seed value and generates a result that approximates a random number. This property allows us to simulate randomness that is present in real world systems. In this fashion, an agent based simulation that provides the same input variables but implements a level of randomness can produce, sometimes, significantly different outcomes. A key challenge of analysing an ABM is to identify an appropriate set of state variables. Markov decision processes is a technique where decision making is partly stochastic and partly under the control of a decision maker (Papadimitriou and Tsitsiklis, 1987; White and White, 1989). A number of multi agent simulations have adopted this approach to modelling an agents decision process (Boutilier, 1999; Becker et al., 2003). This technique introduces randomness and uncertainty into an agents decision. It diverges from the utilitarian approach of perfect agent rationality used in some of the models described above (Centola et al., 2005; Bikhchandani et al., 1998) etc. We know from real life that human being often behave in a manner that is anything but rational (Denes-Raj and Epstein, 1987; Ellickson, 1989).

In this related research section we discussed differing models of norm dissemination and the associated discipline of cultural evolution. We also discovered that norm emergence is heavily influenced by the individuals that an agent meets in the network. Real world interactions are dynamic, this is a feature we aim to capture in this paper. We have also gained an understanding of some of the key metrics that are used for analysing network models. We aim to use this understanding to formulate a model which we describe in the next section.

---

<sup>4</sup> The degree of an agent is the number of acquaintances it has in its social network.

<sup>5</sup> A Monte Carlo algorithm relies on repeated random sampling to compute their results.

**Table 1** Experiments Conducted

	Rational Agents	Markov Agents
Binary Norm	Case 1	Case 2
n-bit Norm	Case 3	Case 4

### 3 Model Design

In the following section we first describe our weighted random interaction algorithm and then formally define the decision making rules agents use when updating their normative beliefs. Table 1 shows the set of experiments that were conducted. They are divided into simulations where agents act completely rationally and where their decision is determined probabilistically using a Markov Decision Process. Furthermore we examine the cases where agents simply have to choose between one of two norms and where an n-bit norm propagates through the population.

#### 3.1 Weighted Random Interactions

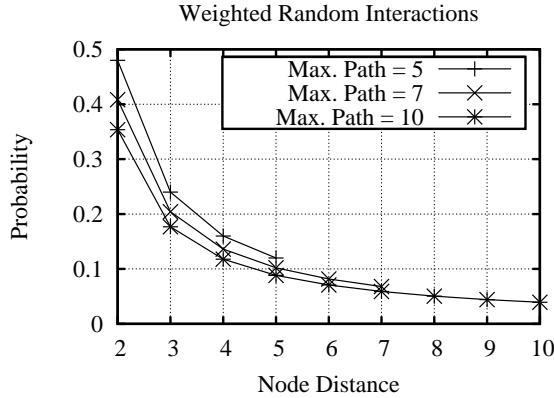
Agents interact with random members of the population using what we describe as a Weighted Random Interaction (WRI) algorithm based on their distance from others on the network. We use a modified version of Zipf's law, Equation 2, to calculate a node's weight. The probability of agent  $i$ , with Maximum Path Distance (MPD) of  $M$ , randomly interacting with agent  $h$  having a path distance of  $d$  from  $i$  is equal to:

$$p_{ih}(d) = \frac{1}{\sum_{m=1}^{M-1} \left( \frac{1}{m^\lambda} \right)} \quad d \geq 2 \quad (2)$$

Where  $\lambda$  is the exponent that characterises the distribution. For simplicity in the experiment carried out in this paper we set  $\lambda = 1$ . Note the condition that  $d \geq 2$  as a node is assumed to interact with members of its social network ( $d = 1$ ). It can be seen from Fig. 2 that the distribution is normalised and the frequencies sum to 1 as expressed in Equation 3.

$$\sum_{d=2}^{M-1} p_{ih}(d) = 1 \quad (3)$$

The graph shown in Fig. 2 shows the distance probability distribution of three different nodes with Maximum Path Distance (MPD) ranging from 5 to 10. We can see from the diagram that agents with a lower path distance are more likely to interact than ones with a higher path distance.



**Fig. 2** Weighted Random Interactions

### 3.2 Case 1 - Rational Agents, Binary Norm

An Agent receives a utility from observing the norms that have been adopted by the other individuals it encounters. Agent  $i$  interacts with  $k_i$  members of their social network and  $r_i$  randomly selected agents. Initially nodes are set to having adopted either social convention  $m$  or  $n$ . In all the models described here from Section 3.2 to Section 3.5 nodes interact with a period drawn randomly from an exponential distribution with mean duration  $\epsilon_i = 3$ . This models the fact that all agents don't update their norm selection simultaneously. An agent,  $i$ , will chose to adopt norm  $m$  if the utility it observes from adopting this norm is greater than the utility it perceives from adopting convention  $n$  as defined in Equation 4.

$$u_{i,t}^m > u_{i,t}^n \quad (4)$$

The utility that agents perceive from each norm is defined in Equation 5. This is divided into the utility communicated from its direct neighbours,  $D_{i(t-1)}^j$ , plus the perceived utility from the random interactions it makes,  $R_{i(t-1)}^j$ .

$$u_{i,t}^m = \alpha D_{i(t-1)}^m + \beta R_{i(t-1)}^m \quad (5)$$

where  $\alpha$  is the weighting placed on an agent interacting with the members of its own social network and  $\beta$  is the weighting of interactions taking place with random members of the agents network. The higher the  $\beta$  value the more importance agents place on random interactions. The direct network effects are defined in Equation 6 where  $n$  is the total number of nodes on the network and  $\theta_{h(t-1)}^m = 1$  if agent  $h$  has adopted social convention  $m$ .

$$D_{i(t-1)}^m = \sum_{h=1}^n \mu_{ih} \theta_{h(t-1)}^m \quad \mu_{ih} \begin{cases} 1 & \text{if } i \text{ is an acquaintance of } h \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

Similarly we define the random network effects in Equation 7 where  $n$  is the total number of nodes on the network and  $\omega_{h(t-1)}^m = 1$  if agent  $h$  has adopted social convention  $m$ .

$$R_{i(t-1)}^m = \sum_{h=1}^n \phi_{ih} \omega_{h(t-1)}^m \quad \phi_{ih} \begin{cases} 1 & \text{if } i \text{ has a random interaction with } h \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

### 3.3 Case 2 - Markov Decision Process Agents, Binary Norm

In this case agents calculate their perceived utility as in Equation 5. But instead of choosing the norm with the highest utility they assign a weighting to each norm and select one using a Markov Decision Process. The more observations an agent has of a norm the greater the chance it will be adopted as expressed in Equation 8.

$$P(i_t^m) = \frac{u_{i,t}^m}{u_{i,t}^m + u_{i,t}^n} \quad (8)$$

Where  $P(i_t^m)$  is the probability of agent  $i$  choosing norm  $m$  at time  $t$  and  $u_{i,t}^m$  and  $u_{i,t}^n$  is the utility that agent  $i$  perceives from norms  $m$  and  $n$  respectively. Case 1 assumed that all agents used perfect rationality when calculating the most efficient norm to adopt whereas this decision making process models the uncertainty and irrationality that an individual agent may have when deciding on which norm to pick.

### 3.4 Case 3 - Rational Agents, n-bit Norm

Next we define the norm as being a set of  $n$  binary bits that describe its characteristics. The total number of unique states that the norm can therefore have is  $2^n$ . We assume that if the characteristics of a norm of node  $i$  are too different from those of node  $h$  then there will be no alteration to node  $i$ 's norm. We calculate the level of similarity between the norms of  $i$  and  $h$  from Equation 9 as:

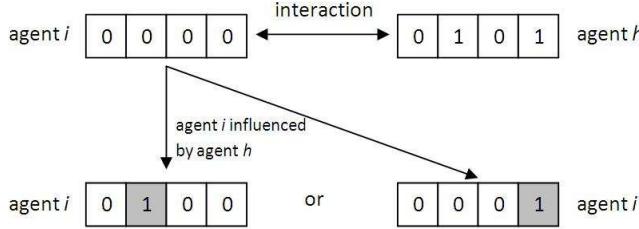
$$\delta_{ih} = \frac{\sum_{p=1}^n \kappa_p^{ih}}{n} \quad \kappa_p^{ih} \begin{cases} 1 & \text{if } i \text{ and } h \text{ share the same characteristic} \\ & \text{at bit position } p \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

If  $\delta_{ih} \geq d$  then choose a random bit position,  $p$ , such that  $i_p \neq h_p$  and set  $i_p = h_p$  where  $d$  is the bounded confidence tolerance parameter and  $i_p$  and  $h_p$  is the value of the bit at position  $p$  of norm for agents  $i$  and  $h$  respectively. We can set  $d$  in the range of  $[0 : 1]$ . The higher the value of  $d$  the more similarity there needs to be between norms before they will have an effect on one another. For example, if agent  $i$  has a norm that is characterised by the 4-bit sequence

0000 and agent  $h$  has a norm defined by the sequence 0101 then the similarity between the two norms is  $\delta_{ij} = 0.5$  as shown in Fig. 3. If  $d = 0.25$  then agent  $h$ 's norm will influence agent  $i$ 's and a single bit will be changed to bring agent  $i$ 's norm closer to agent  $h$ 's. At each decision time step agents interact with one member from either their social network or the random acquaintances they meet. As in Case 1 and Case 2 above, agents interact with a period drawn randomly from an exponential distribution with mean duration  $\epsilon_i = 3$ . An agents' choice of who to interact with is a function of both the strength and number of both social and random contacts. Therefore, the probability that an agent will interact with a member of its own social network,  $P(i_s)$ , is defined in Equation 10.

$$P(i_s) = \frac{\alpha k_i}{(\alpha k_i) + (\beta r_i)} \quad (10)$$

where  $k_i$  is the size of agents  $i$ 's social network and  $r_i$  is the number of random contacts that agent  $i$  has. Similarly, the probability that agents  $i$  will interact with a random agent not in its social network is defined as  $P(i_r) = 1 - P(i_s)$ . We can see that the greater number of random contacts that an agent has the greater chance they have of interacting and being influenced by these random contacts. Note that the randomness in the model stems from which bit position to change rather than whether to change in the first place. In this regard an agent takes a utilitarian or completely rational approach each time when deciding whether to change their norm or not.



**Fig. 3** norm influence on agent  $i$  by agent  $h$

### 3.5 Case 4 - Markov Decision Process Agents, n-bit Norm

In Case 4 agents now decide whether to alter their n-bit binary norm probabilistically based on the similarity of the norms as described by (Axelrod, 1997; Kuperman, 2006; Centola *et al.*, 2007). At each decision time step an agent will choose a member of either its social network or random contacts to interact with as defined in Equation 9. From Equation 9 we can see that  $\delta_{ih}$

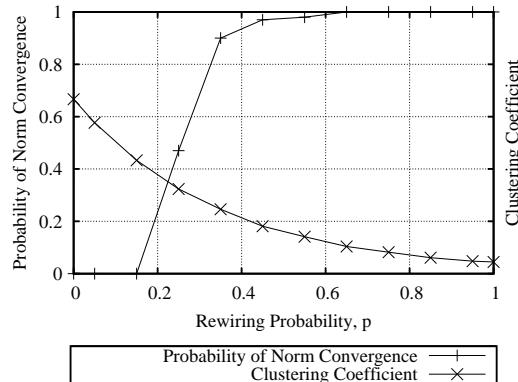
specifies the level of similarity between the two norms. Agent  $i$  will update its own norm with probability  $P(i) = \delta_{ih}$ . The more similarity that agents have the more likely it is that their norms will influence one another. We can see from this approach that an agent may be influenced to alter its norm even if it observes a norm that is very different from its own.

## 4 Results

In this Section we will outline the results we obtained from the models described in Section 3. In Section 4.1 we examine the effects of norm convergence when simply changing the rewiring probability of a SW Network. For the rest of the experiments we generate a fixed SW Network with a population of 1000 agents with average degree of  $k_i = 10$  and  $p_w = 0.05$ . We then add random interactions to the simulation using the four different scenarios described in Table 1. All the results shown are the average of 100 different simulations unless otherwise stated. A new SW network was generated for each simulation.

### 4.1 Varying Rewiring Probability

Fig. 4 shows the effect of norm convergence when the rewiring probability is changed. We observe that the probability of all agents converging on a common norm is increased when the value of  $p_w$  is increased. This conforms to the finding of Lee et al. (2006) in the domain of product adoption mentioned earlier. While increasing the value of  $p_w$  results in an increase in norm convergence, it reduces the level of clustering in the network. Real world human networks have high levels of clustering so this means that increasing  $p_w$  is unrealistic. In this section we have learnt:



**Fig. 4** Probability of Norm Convergence varying  $p$

**Table 2** Simulation Experiment Variables

N	$\alpha$	$p_w$	$k_i$	$\epsilon_i$	Number of Norms
1000	1.0	0.05	10	3	2

- Norm convergence won't occur on small world networks with low  $p_w$
- Exceeding a threshold level of  $p_w$  results in convergence in all cases.
- Increasing  $p_w$  reduces the level of clustering on the network

In the next experiments we maintain a core, highly clustered, small world network but introduce ad hoc random interactions that the agents have with others in the population.

#### 4.2 Random Interactions

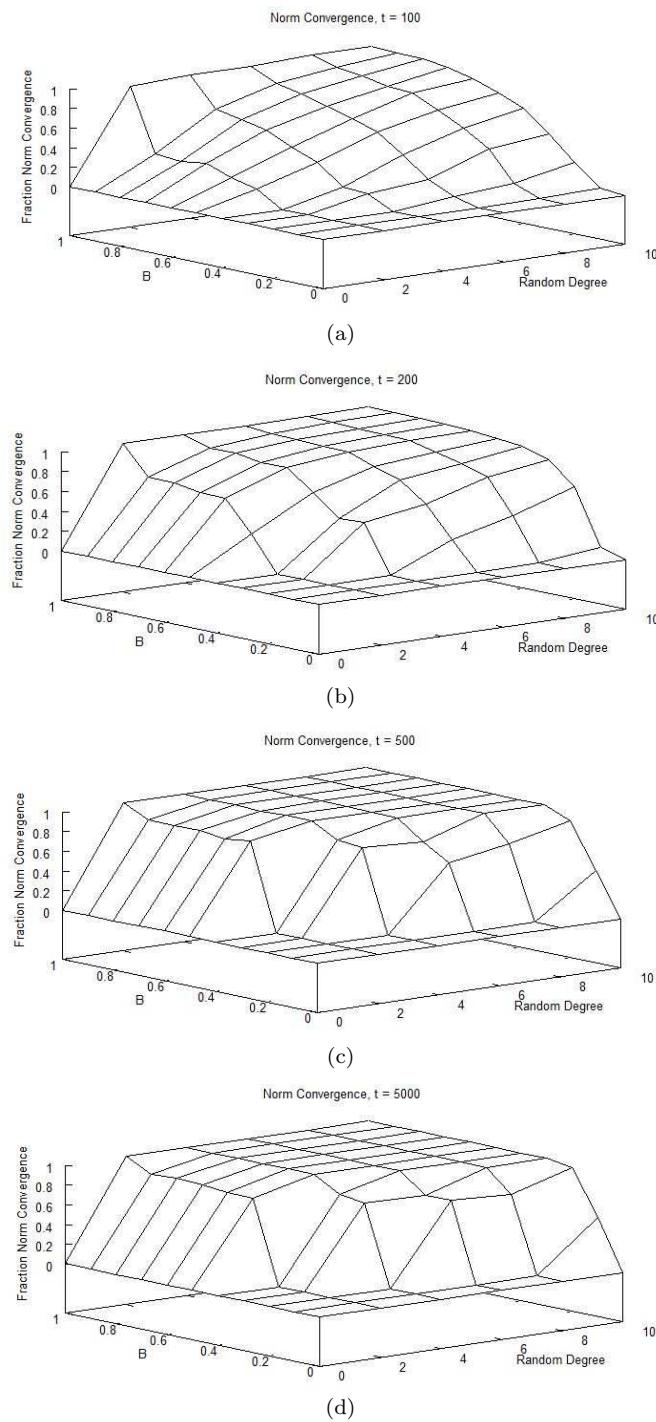
In the following experiments, Sections 4.2.1 - Sections 4.2.4, a SW network is created with a rewiring probability of  $p_w = 0.05$  and  $\alpha = 1$  as shown in Table 2. As we have seen in Fig. 4, norm convergence will not happen when  $p_w$  is at this level.

Initially each agent on the network maps its social distance from every other agent. We used Dijkstra's algorithm (Johnson, 1973) to calculate the MPL for each node as described in Section 3.1. Every time an agent interacts it generates a new set of ad hoc random interactions based on the WRI algorithm described and discards the old ones.

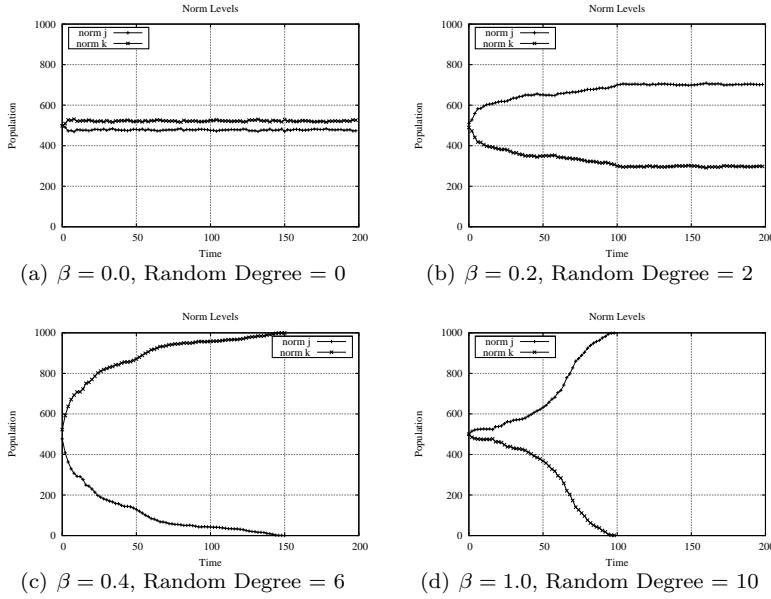
##### 4.2.1 Case 1

In this model agents act in a perfectly rational fashion and choose between either of two competing norms as outlined in Section 3.2. In Fig. 5 we vary both the value agents place on random interactions,  $\beta$ , and the number of random interactions that they have,  $r_i$ . The number of random interactions starts at 0 and is increased by a value of 2 until it reaches 10. The strength of random interactions,  $\beta$ , is increased from 0 to 1. When  $\beta = 1$  and  $r_i = 10$  then an agent has the same number of random interaction with members of the population as their social network and places the same strength on these random interactions.  $\beta$  therefore is a measure of how much importance agents place on encounters with individuals who are not in their social network. A low  $\beta$  could represent an agent passively observing a norm whereas a high value could represent an active conversation etc.

Fig. 5 shows a snap shot of the level of norm convergence over a series of different time steps. In Fig. 5(a) when  $t = 100$  the level of norm convergence is high when both  $\beta$  and the number of random interaction are high. Norm convergence fails to occur when both the strength and quantity of random interactions is too low. Indeed, when agents are having up to 8 random interactions but those occurrences only carry a weight of 0.1 of random interactions then norm emergence will not occur. Increasing the number of random



**Fig. 5** Level of Norm Convergence over Time



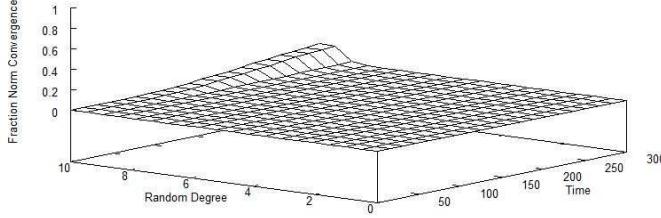
**Fig. 6** Sample Run over Time, Agents acting Rationally

interaction or increasing the strength of these interactions results in norm emergence. From Fig. 5 we can see that if  $\beta > 0.5$  and the number of random interactions  $r > 6$  then norm convergence is guaranteed. If agents have the same number of random interactions as members of their social network, or  $k_i = r_i = 10$ , then  $\beta$  only needs to be 0.2 to guarantee norm convergence. Interestingly, allowing agents to interact beyond  $\sim 300$  time steps has little impact on the level of norm convergence. This experiment shows that random interactions with members of a nodes social network plays an important role in norm convergence.

Fig. 5 showed the end result of an average of 100 simulations. Next, in Fig. 6, we look at a sample run with varying levels of  $\beta$  and  $r_i$ . We can see from Fig. 6(a) that when agents are just interacting with their own social network then the agents quickly settle on two large groups of norms of approximately half the population each. Increasing both the level and strength of random interactions results in a larger portion of the population adopting the same norm as in Fig. 6(b). The random interactions are not enough to overcome the local norm bias'. Increasing the strength and number of random interactions further will results in norm convergence in quicker times as Fig. 6(c) and Fig. 6(d) show.

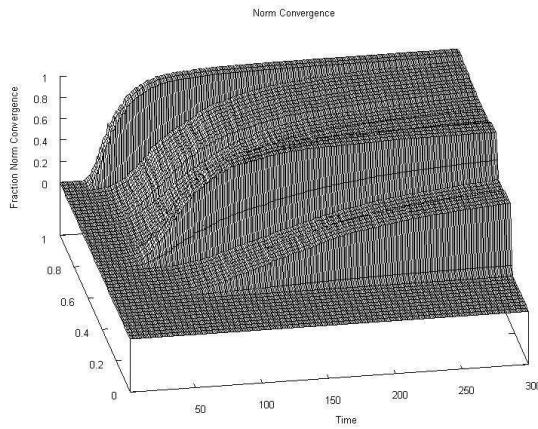
In Fig. 7 we set the level of  $\beta = 0.4$  and increment the number of random interactions. We can see that there is no norm convergence when the number of random interactions is low. Once members of the population have enough random interactions then there is a steady increase in the number of simula-

tions resulting in convergence. The graph appears like a terrace because the random degree is increased in steps of 2.



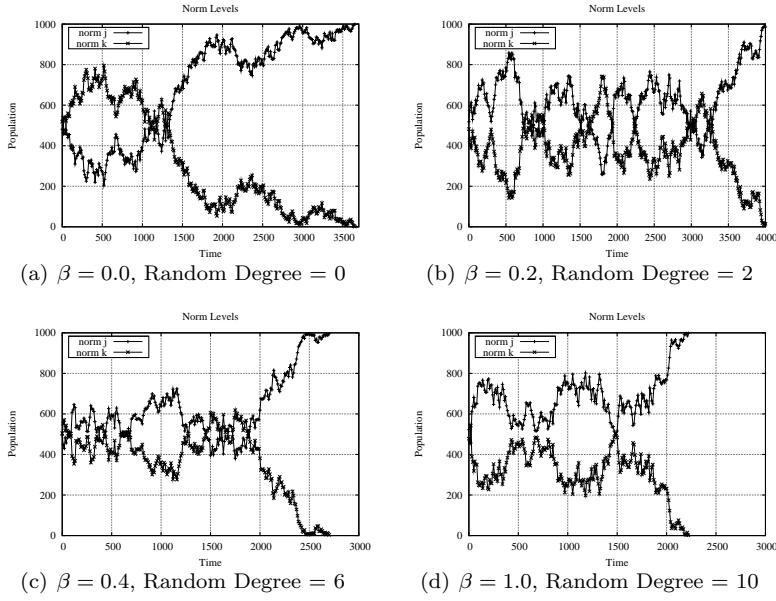
**Fig. 7** Level of Norm Convergence over Time Varying Random Degree

In Fig. 8 we set the level of random interactions to 4 and incremented the level of  $\beta$ . We can see that there are several jumps in norm convergence when we increase the level of  $\beta$ . Specifically, when  $\beta \leq 0.2$  then none of the simulations converge to a common norm. When  $\beta \geq 0.5$  then all the simulations converge to a common norm. We can also see that when  $0.2 \leq \beta \leq 0.3$  then some norm emergence does occur but at a much slower rate.



**Fig. 8** Level of Norm Convergence over Time Varying Beta

In all the simulations carried out in this section we can see that unless the number and strength of random interactions is enough the population fails to converge to a single norm. This is because agents create local group of reenforcing norms. It takes random encounters from outside this group to break their local biases. In summary, the experiments on Case 1 as outlined in this section we find that:



**Fig. 9** Sample Run over Time, Agents acting Probabilistically

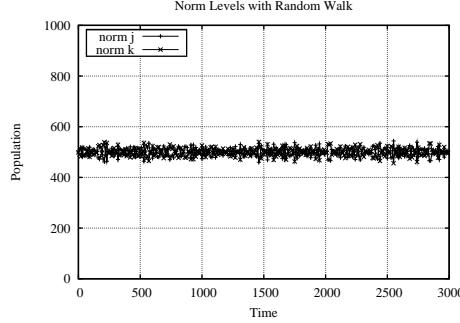
- Norm convergence won't occur if there is not sufficient random interactions.
- An increased number of random interactions results in convergence at a faster rate.
- When agents act with perfect rationality then equilibrium states can exist with multiple norms present.

#### 4.2.2 Case 2

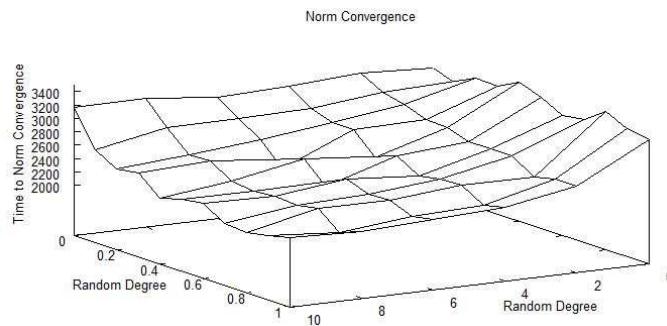
Next agents decide on which of two competing norms to adopt based on a Markov Decision Process as defined in Section 3.3. In this case the agents will converge on a single norm in *every* simulation. This is because the stochastic behaviour of the agents will result in local biases being broken within the population. Fig. 9 shows a sample of one run for varying levels of  $\beta$  and  $r_i$ . The volatile nature of the norms is clear from this graphs. When agents appear to be converging on a global consensus there can be a sudden shift in the opposite direction. These shifts are due to the sometimes irrational decision making that we have now introduced into agent behaviour. If we compare this to Fig. 10 which shows the level of norm convergence if agents were following a random walk. Fig. 9 displays a general trend toward norm convergence.

Fig. 9 clearly shows the potential of a norm's popularity to change over time if we assume that humans don't behave as perfectly rational utilitarian decision makers. Fig. 11 shows the average amount of time it takes agents to reach a consensus on which norm to adopt. The rough surface of the graph highlights how randomness in the system effects the outcome. We can clearly

see, however, a general decline in the time required for norm convergence on the network as  $\beta$  and  $r_i$  are increased. The greatest reduction is when  $\beta$  and  $r_i$  are both small which suggests that there is diminishing returns on the network when these variable are continually incremented.



**Fig. 10** Norm Levels under Random Walk

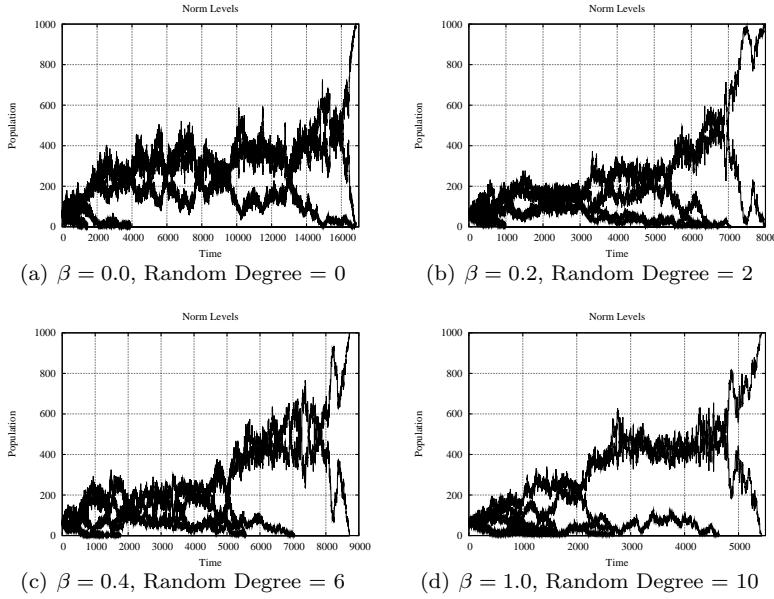


**Fig. 11** Time to reach Norm Convergence

This section has highlighted a number of aspect of norm convergence, specifically we find that:

- Agents that behave with non-perfect rationality results in norm convergence in every simulation.
- Increasing the number of random interactions that agents have results in a reduction in the time required for norm convergence.
- This reduction in time is more prominent when  $\beta$  and  $r_i$  are small, with diminishing returns thereafter.

In the next section we look at the effect of norm *evolution* when agents are interacting with random members of the population.



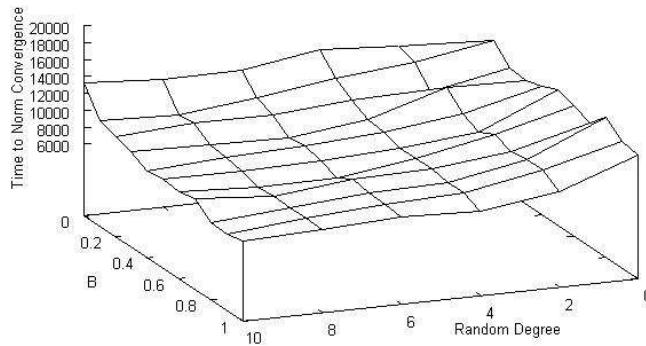
**Fig. 12** Sample Run over Time, Agents acting Rationally.  $d = 0.25$

#### 4.2.3 Case 3

In this case if agents observe a norm that is close enough to their own they will alter their own  $n$ -bit norm to become closer to the one they observe as described in Section 3.4. Once the norm reaches a similarity level that exceeds a certain threshold level then an agent will be influenced by its observation. In both Case 3 and in Case 4 below we set  $n = 4$ . There can therefore be  $2^4 = 16$  unique representations of the norm. In many of the papers discussed in Section 2.1 on cultural evolution the value of  $n$  is altered. As we are primarily concerned with the effect of random interactions here we leave this for future work. As described in Section 3.4 the tolerance parameter,  $d$ , defines how similar the norms must be in order to interact with one another. Fig. 12 shows a sample of four runs with increasing levels of both  $\beta$  and  $r_i$  when the tolerance parameter is set to  $d = 0.25$ . Each line on the graphs represent the level  $a$  of single  $n$ -bit norm. In this case agent  $i$  only needs to share a single bit in common with agent  $k$  during an interaction in order to be influenced by it. We find that when the level of  $d$  is this low then agents converge to a single norm in every simulation. In this case the cognitive model we have constructed for agents implies that they are easily influenced by others in the population even if the norm they observe has more differences than similarities. Notice the pattern that emerges from each of the four graphs displayed in Fig. 12. In each case two very similar norms spread throughout the population and then diverge. In Fig. 12(b) we can see that at around time  $t = 5500$  that

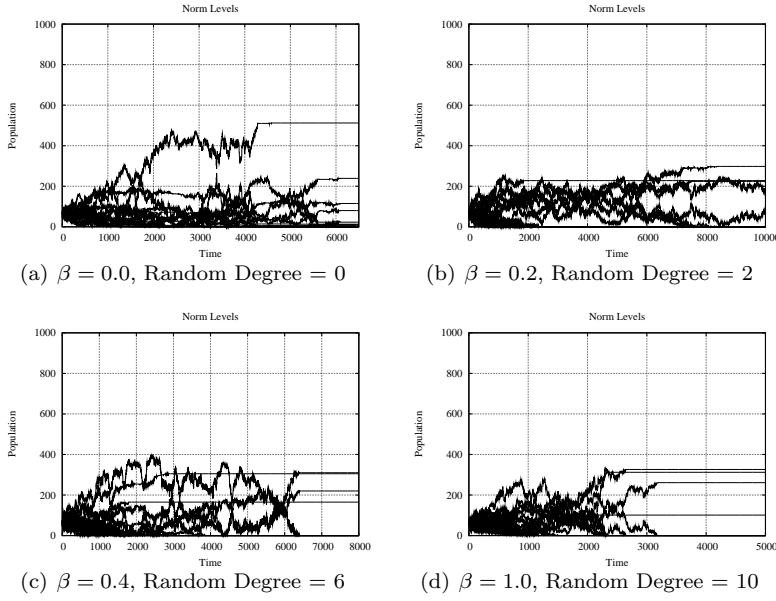
two norms start to be generally adopted by population. These two norms are represented by the bit strings 0001 and 1001. The two norms only have one bit difference from one another. Once the population has decided that this is the general form that the norm will take all other versions of the norm die out. Then finally at around time  $t = 7000$  the population diverges and settles on a final norm. In Fig. 12(d) the two norms 0010 and 1010 dominate much of the simulation. Once other representations of the norm have been eliminated then the population settles on a final social norm.

Looking at the amount of time required for agents to converge on a norm we see in Fig. 13 that there is an initial reduction in convergence time when the values of  $\beta$  and  $r_i$  are increased but this levels off once the values incremented further. As agents are only interacting with one member of either their social network or random contacts we can see that an increase in the number of  $r_i$ , while increasing the likelihood of interacting with a random contact, only changes the spread of possible random contacts. And as random contacts are weighted towards agents with lower social distance on the network an increase in  $r_i$  increases the chance of an agent interacting with an individual with very high social distance.



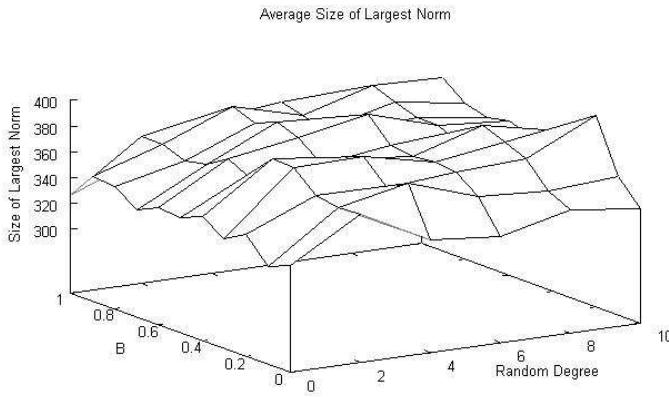
**Fig. 13** Time to reach Norm Convergence,  $d = 0.25$

Next we increased the tolerance parameter to  $d = 0.75$ . In this case agents need to share three out of four bits of their norm in order to be influenced by their contacts. We can see from Fig. 14 that a form of polarisation occurs whereby agents fall into segregated groups of norms that are not influenced by others that they meet. As agents in this scenario require a large amount of similarity before they will be influenced by others they are unwilling to change their norm. Some of the groups can be quite large and represent a large fraction of the total population as in Fig. 12(a). Other times it can splinter into a series of smaller norms all co-existing on the network. We discovered that as the value of  $\beta$  and  $r_i$  are increased agents settle on a final norm choice faster. Next we look at the effect that changing the value of  $\beta$  and  $r_i$  had on the final size of



**Fig. 14** Sample Run over Time, Agents acting Rationally.  $d = 0.75$

norms in the population. To that end we measured the size of the largest group in the population once the agents had all settled on a specific norm. Fig. 15 shows the size of the largest component on the network. Notice how there is a slight increase in the size of the largest component when we initially increase  $\beta$  and  $r_i$ . The rough surface of Fig. 15 demonstrates that even through these results are the averages of 100 simulations there is still much variability in the final outcome.



**Fig. 15** Size of Largest Component when  $d = 0.75$

These results indicate that increasing the value of  $\beta$  and  $r_i$  has a less pronounced effect on the development and dissipation of the norm when the norms themselves are considered transient constructs that are capable of changing under constrained agent interactions. In this section we have found that:

- Agents will converge to a common norm provided the tolerance parameter is low enough.
- If the tolerance parameter is too high then polarisation into different norm subgroups will occur.
- Increased levels of random interactions fails to have a significant effect on the size of these subgroups.

#### 4.2.4 Case 4

The following section shows the effect of norm evolution when the propensity of an agent to alter its norm is proportional to the similarity of the norms. Unlike Case 3 above, agents in this model may not act in the same fashion each time given the same information. This adds an element of cognitive agent randomness into the decision making process<sup>6</sup> as in Case 2 above. We can see from Fig. 16 that the general behaviour of agents in the population is similar to Case 3 above when  $d = 0.25$ . That is, two similar versions of the norm rise above the rest and then finally diverge. This indicates that the population is deciding on the general form of the norm before finally converging on the exact 4-bit characteristic that is adopted by the whole population.

As every simulation results in convergence to a single norm we are interested in investigating how long it takes the population to converge when the values of  $\beta$  and  $r_i$  are increased. Fig. 17 shows the fraction of simulations that have converged to a single norm at different time steps. Fig. 17(a) shows that there is initially an increase in the level of norm convergence when  $\beta$  and  $r_i$  are increased. Beyond this the increase is less prominent and given enough time all of the simulations converge in a single norm as in both Fig. 17(b) and Fig. 17(c).

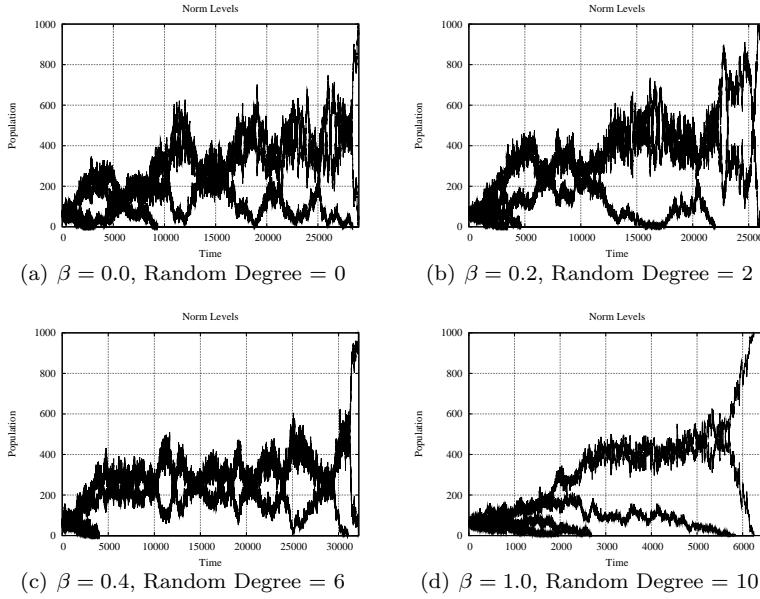
In this section, when agents act in a non-rational manner, our experiments indicate that:

- Random interactions are less of an influence on norm evolution compared to binary norm dissipation.
- Agents will modify their norm to suit others rather than completely abandoning it as in the binary case.
- Increasing the number of random interaction doesn't significantly effect the speed at which a consensus is formed.

Section 4 has shown us the importance of clearly defining both the cognitive framework of the agents and our definition of both what a norm is and how it spreads. In the first scenario agents choose between two competing binary

---

<sup>6</sup> If two norms are completely opposite from one another such as 0101 and 1010 then their possibility of effecting one another is zero as  $\delta_{ih} = 0$ .

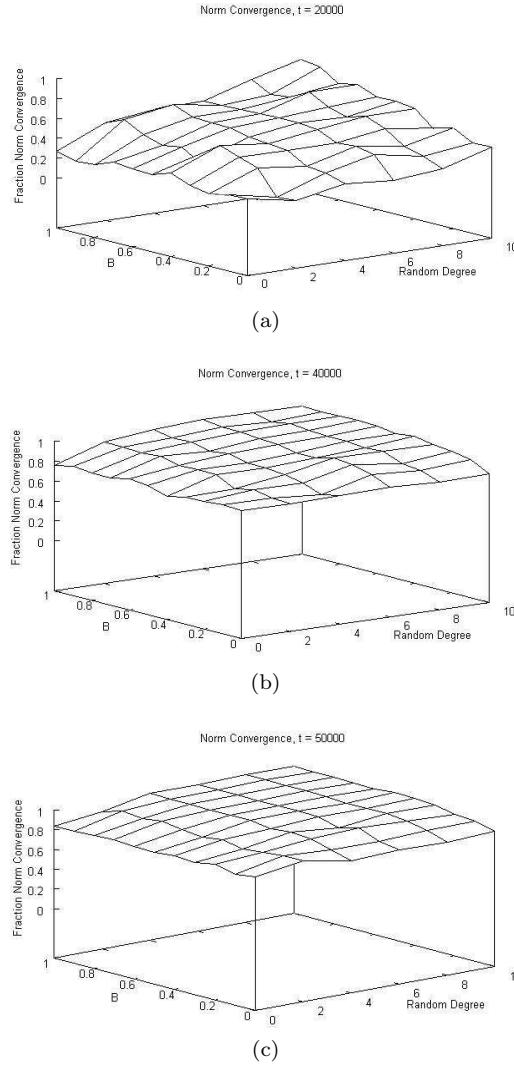


**Fig. 16** Sample Run over Time, Agents acting Probabilistically

alternatives. We observed a clear increase in norm convergence and in a quicker time when the number of random interactions is increased. When the norm can change over time based on an agents' influence we don't see the same dramatic convergence effect from random interactions. Rather, the propensity an agent has to be influenced by its interactions plays an important role.

## 5 Conclusions

The aim of this paper was to construct a more realistic network of agent interactions that can generate insight into the emergence of norms in society. We defined an algorithm that uses a nodes social distance on the network to calculate its chance of interacting with a random member of the population. We believe that this method of weighting random ad hoc interactions based on an agents social distance on the network represents a novel way of modelling real human interactions. We then investigated norm dissemination on populations of agents interaction on small world networks with random contacts. We separated our experimentation into categories that defined both the cognitive behaviour of the agents and the characteristics of the norm that was under investigation. We demonstrated the importance of random interactions in populations where agents had to choose between two competing norms. If agents behaved in a perfectly utilitarian manner then random interactions were vital in breaking local biases and allowing the population to converge to a single norm. When an agents cognitive behaviour was defined as a Markov Decision



**Fig. 17** Level of Norm Convergence over Time. Agents acting Probabilistically

process then random interaction resulted in a reduction in the amount of time required for norm convergence. When we defined the norm as an n-bit sequence that characterised its behaviour we discovered that increased random interactions had a less pronounced effect on the population. This is because increased random interactions resulted in the norm itself changing over time thus negating some of the benefits from accessing agents outside ones own social network. Increasing the amount of similarities between norms required

---

for influence resulted in a splintering effect where different groups form all adopting a separated norm.

This research has highlighted the importance of clearly defining both the cognitive decision making process of agents and the norm representation under investigation. Our simulations produced a very different set of outcomes under alteration of these conditions.

There were several areas of this research that will be explored in the future. As we were mainly focusing on the effect of random interactions we only tested two values for the tolerance parameter  $d$ . We also hope to run additional sensitivity analysis on other variable involved in the simulation. We hope to expand this work by giving agents a memory of previous interactions. How would this effect the dissemination of norms? We could also address the issue of norm violation and punishment which we hope to pursue going forward.

## References

- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *The Journal of Conflict Resolution*, **41**(2), 203–226.
- Barrat, B. K. . A. (2008). Consensus formation on adaptive networks. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, **77**(1), 016102.
- Baum, J. A. C., Shipilov, A. V., and Rowley, T. J. (2003). Where do small worlds come from? *Ind Corp Change*, **12**(4), 697–725.
- Becker, R., Zilberstein, S., Lesser, V., and Goldman, C. V. (2003). Transition-independent decentralized markov decision processes. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 41–48, New York, NY, USA. ACM.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, **100**(5), 992.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1998). Learning from the behavior of others: Conformity, fads, and informational cascades. *The Journal of Economic Perspectives*, **12**(3), 151–170.
- Boutilier, C. (1999). Sequential optimality and coordination in multiagent systems. In *IJCAI '99: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 478–485, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Carley, K. M. (2009). Computational modeling for reasoning about the social behavior of humans. *Computational & Mathematical Organization Theory*, **15**(1), 47–59.
- Centola, D., Willer, R., and Macy, M. (2005). The emperor's dilemma: A computational model of self-enforcing norms. *American Journal of Sociology*, **110**(4), 1009–1040.

- 
- Centola, D., Gonzalez-Avella, J. C., Eguiluz, V. M., and San Miguel, M. (2007). Homophily, Cultural Drift, and the Co-Evolution of Cultural Groups. *Journal of Conflict Resolution*, **51**(6), 905–929.
- Conte, R., Falcone, R., and Sartor, G. (1999). Introduction: Agents and norms: How to fill the gap? *Artificial Intelligence and Law*, **7**(1), 1–15.
- Davis, G. F., Yoo, M., and Baker, W. E. (2003). The Small World of the American Corporate Elite, 1982–2001. *Strategic Organization*, **1**(3), 301–326.
- Deffuant, G., N. D. A. F. and Weisbuch, G. (2001). Mixing beliefs among interacting agents. *Advances in Complex Systems*, (3), 87–98.
- Denes-Raj, V. and Epstein, S. (1987). Conflict between intuitive and rational processing: When people behave against their better judgment. *Journal of Personality and Social Psychology*, **66**(5), 819.
- Ehrlich, P. R. and Levin, S. A. (2005). The evolution of norms. *PLoS Biol*, **3**(6), e194.
- Ellickson, R. (1989). Bringing culture and human frailty to rational actors: A critique of classical law and economics. *65 CHI-KENT L. REV.*, **23**.
- Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, **11**(4), 12.
- Erdos, P. and Reyni, A. (1961). On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, **4**(5), 17–61.
- Fenner, T., Levene, M., Loizou, G., and Roussos, G. (2007). A stochastic evolutionary growth model for social networks. *Comput. Netw.*, **51**(16), 4586–4595.
- Fronczak, A., Fronczak, P., and Holyst, J. A. (2002). Average path length in random networks.
- Giddens, A. (1984). *The constitution of society : outline of the theory of structuration / Anthony Giddens*. Polity, Cambridge ..
- Goldfarb, B. and Henrekson, M. (2003). Bottom-up versus top-down policies towards the commercialization of university intellectual property. *Research Policy*, **32**(4), 639 – 658.
- Gonzalez-Avella, J. C., Cosenza, M. G., and Tucci, K. (2005). Nonequilibrium transition induced by mass media in a model for social influence.
- Henrick, J. and Boyd, R. (2001). Why people punish defectors: Weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *Journal of Theoretical Biology*, **208**(1), 79 – 89.
- Horne, C. (2007). Explaining Norm Enforcement. *Rationality and Society*, **19**(2), 139–170.
- Izquierdo, L. R., Izquierdo, S. S., Galán, J. M., and Santos, J. I. (2009). Techniques to understand computer simulations: Markov chain analysis. *Journal of Artificial Societies and Social Simulation*, **12**(1), 6.
- Johnson, D. B. (1973). A note on dijkstra's shortest path algorithm. *J. ACM*, **20**(3), 385–388.
- Kittock, J. (1995). Emergent conventions and the structure of multi-agent systems. In *Lectures in Complex systems: the proceedings of the 1993 Complex systems summer school, Santa Fe Institute Studies in the Sciences of*

- 
- Complexity Lecture Volume VI, Santa Fe Institute*, pages 507–521. Addison-Wesley.
- Klein, R. G. (1999). *The Human Career: Human Biological and Cultural Origins*. University Of Chicago Press.
- Kuperman, M. N. (2006). Cultural propagation on social networks. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, **73**(4), 046139.
- Lee, E., Lee, J., and Lee, J. (2006). Reconsideration of the Winner-Take-All Hypothesis: Complex Networks and Local Bias. *Management Science*, **52**(12), 1838–1848.
- Liefbroer, A. C. and Billari, F. C. (2009). Bringing norms back in: a theoretical and empirical discussion of their importance for understanding demographic behaviour. *Population, Space and Place*, pages 1544–8452.
- Lopez-Pintado, D. and Watts, D. J. (2008). Social Influence, Binary Decisions and Collective Dynamics. *Rationality and Society*, **20**(4), 399–443.
- Lopez y Lopez, F., Luck, M., and d'Inverno, M. (2006). A normative framework for agent-based systems. *Computational & Mathematical Organization Theory*, **12**(2), 227–250.
- McKeown, G. and Sheehy, N. (2006). Mass media and polarisation processes in the bounded confidence model of opinion dynamics. *Journal of Artificial Societies and Social Simulation*, **9**.
- McPherson, M., Smith-Lovin, L., and Cook, J. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, **27**, 415–444.
- Milgram, S. (1967). The small world. *Psychology Today*, **2**, 6067.
- Mukherjee, P., Sen, S., and Airiau, S. (2008). Norm emergence under constrained interactions in diverse societies. In *AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 779–786, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.
- Newman, M. E. J. (2002). Random graphs as models of networks.
- Ostrom, E. (2000). Collective action and the evolution of social norms. *The Journal of Economic Perspectives*, **14**(3), 137–158.
- Papadimitriou, C. H. and Tsitsiklis, J. N. (1987). The complexity of markov decision processes. *Mathematics of Operations Research*, **12**(3), 441–450.
- Rahmandad, H. and Sterman, J. (2008). Heterogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models. *Management Science*, **54**(5), 998–1014.
- Ruef, M., Aldrich, H. E., and Carter, N. M. (2003). The structure of founding teams: Homophily, strong ties, and isolation among u.s. entrepreneurs. *American Sociological Review*, **68**(2), 195–222.
- Savarimuthu, B. T. R., Cranefield, S., Purvis, M., and Purvis, M. (2007). Norm emergence in agent societies formed by dynamically changing networks. In *IAT '07: Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 464–470, Washington, DC, USA. IEEE Computer Society.
- Shilling, C. (1999). Towards an embodied understanding of the structure/agency relationship. *The British Journal of Sociology*, **50**(4), 543–562.

- 
- Shoham, Y. and Tennenholtz, M. (1995). On social laws for artificial agent societies: Off-line design. *Artificial Intelligence*, **73**, 231–252.
- Verspagen, B. and Duysters, G. (2004). The small worlds of strategic technology alliances. *Technovation*, **24**(7), 563 – 571.
- Villatoro, D., Malone, N., and Sen, S. (2009). Effects of interaction history and network topology on rate of convention emergence. *Proceedings of 3rd International Workshop on Emergent Intelligence on Networked Agents*.
- Walker, A. and Wooldridge, M. (1995). Understanding the emergence of conventions in multi agent systems. *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS '95)*, (1), 384–389.
- Watts, D. (1999a). Small worlds: The dynamics of networks between order and randomness. In *Princeton University Press*.
- Watts, D. and Strogatz, S. (1998). Collective dynamics of 'small-world' networks. *Nature*, pages 440–442.
- Watts, D. J. (1999b). Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology*, **105**(2), 493–527.
- Watts, D. J. and Dodds, P. S. (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research*, **34**(4), 441–458.
- White, C. C. and White, D. J. (1989). Markov decision processes. *European Journal of Operational Research*, **39**(1), 1–16.
- Zhan, F. B. and Noon, C. E. (1998). Shortest Path Algorithms: An Evaluation Using Real Road Networks. *Transportation Science*, **32**(1), 65–73.