

Establishing Norms for Network Topologies

Samhar Mahmoud¹, Nathan Griffiths², Jeroen Keppens¹, and Michael Luck¹

¹ Department of Informatics
King's College London
London WC2R 2LS, UK.
`samhar.mahmoud@kcl.ac.uk`

² Department of Computer Science
University of Warwick
Coventry CV4 7AL, UK.

Abstract. In order to establish a norm in a society of agents, metanorms have previously been proposed as a means of ensuring not that norms are complied with, but that they are enforced. Yet while experimental results have shown that metanorms are effective in fully-connected environments such as that used by Axelrod, there has been limited consideration of such metanorm models with different but more realistic topological configurations. In this paper, therefore, we consider the use of metanorms in supporting norm establishment in lattices and small world networks. Our results suggest that norm establishment is achievable in lattices and small worlds.

1 Introduction

In peer-to-peer systems, agents share resources (hardware, software or information) with others, but if there is no cost to access files nor any limit on the number of files accessible, then there is no incentive to respond to requests nor, more generally, to establish cooperation in the system. Yet cooperation is needed: when self-interested autonomous agents must exchange information without any central control, non-compliance (due to selfish interests) can compromise the entire system. The use of *norms* to provide a means of ensuring cooperative behaviour has been proposed by many [3, 5, 6, 10, 13–15, 17] but, as shown by Axelrod [1], norms alone may not lead to the desired outcomes. In consequence, *metanorms* have been proposed as a means of ensuring not that norms are complied with, but that they are enforced. While experiments have shown that metanorms are effective in fully-connected environments as used by Axelrod, there has been limited consideration of metanorms with different but more realistic topological configurations, which fundamentally change the mechanisms required to establish cooperation.

Some work has already been undertaken on examining the impact of different topologies on norm establishment. For example, Savarimuthu et al. [9] consider the *ultimatum game* in the context of providing advice to agents on whether to change their norms in order to enhance performance for random and scale-free networks. Delgado et al. [4] study norm emergence in coordination games

in scale-free networks, and Sen et al. [11] examine rings and scale-free networks in a related context. Additionally, Villatoro et al. [14] explore norm emergence with memory-based agents in lattices and scale-free networks.

While these efforts provide valuable and useful results, the context of application has been limited, with only two agents involved in each encounter, rather than a larger population of agents. This simplifies the problem when compared with those in which the actions of multiple interacting agents can impact on norm establishment. In particular, Axelrod’s seminal model [1] has provided the foundation for several investigations into norm emergence, yet offers a very general framework, comprising the use of norms and metanorms in populations of agents where the overall behaviour determines whether a norm is established. In this paper we extend Axelrod’s model to address the context of different topological configurations.

The paper begins with an outline of Axelrod’s metanorms game, adjusted to suit the purposes of this paper, and augmented with a learning mechanism. Section 3 then considers the problems that arise from the use of different topologies, and Sections 4 and 5 describe in detail the impact of applying the model in lattices and small worlds.

2 The Metanorms Game

Our model aims to simulate a realistic distributed system in which a community of self-interested agents is encouraged, without being instructed to do so by a central authority, to adhere to a behavioural constraint, or *norm*, that benefits the community but not the individual agent adhering to the norm. This simulation provides an experimental setting that enables us to test under what conditions a situation arises in which the norm governs the behaviour of individual agents.

2.1 Axelrod’s Model

Inspired by Axelrod’s model [1], our simulation focusses only on the essential features of the problem. In the simulation, the agents play a game iteratively; in each iteration, they make a number of binary decisions. First, each agent decides whether to comply with the norm or to defect. Defection brings a reward for the defecting agent, and a penalty to all other agents, but each defector risks being observed by the other agents and punished as a result. These other agents thus decide whether to punish agents that were observed defecting, with a low penalty for the punisher and a high penalty for the punished agent. Agents that do not punish those observed defecting risk being observed themselves, and potentially incur metapunishment. Thus, finally, each agent decides whether to metapunish agents observed to spare defecting agents. Again, metapunishment comes at a high penalty for the punished agent and a low penalty for the punisher.

The behaviour of agents in each round of the game is random, but governed by three variables: the probability of being seen S , boldness B , and vengefulness V . Each round agents are given a fixed number of opportunities o to defect or

Algorithm 1 The Simulation Control Loop: $simulation(T, H, P, E, \gamma, \delta)$

1. **for** each round **do**
 2. interact(T, H, P, E)
 3. learn(γ, δ)
-

comply, each of which has a randomly selected probability of a defection being seen. Boldness determines the probability that an agent defects, such that if an agent's boldness exceeds the probability of a defection being seen then the agent defects. Vengefulness is the probability that an agent punishes or metapunishes another agent. Thus the boldness and vengefulness of an agent are said to comprise that agent's strategy. After several rounds of the game, each agent's rewards and penalties are tallied, and successful and unsuccessful strategies are identified. By comparing themselves to other agents on this basis, the strategies of poorly performing agents are revised such that features of successful strategies are more likely to be retained than those of unsuccessful ones. We need not be concerned with the details of the learning algorithm in this paper, beyond the fact that boldness and vengefulness are simply revised upward or downward as appropriate, in line with a specified learning rate. If most agents employ a strategy of low boldness and high vengefulness, it can be argued that the norm has become *established* in that community, because strategies that lead to defection or to sparing defecting agents are unlikely and lead to high penalties.

2.2 Our Simulation Algorithm

Given Axelrod's model as a starting point, we have previously developed refinements of it that are better suited to real-world distributed systems, by not requiring agents to have information on the private strategies of others, and by allowing agents to improve performance, via a reinforcement learning technique. Since this is not the focus of this paper, we will not provide a full explanation; the full details of why and how are provided in a sister paper [8]. Nevertheless, since these refinements are the starting point for our work here, in this section we briefly review the presentation in [8] to set up subsequent sections.

First, in order to determine the unique effect of each individual action on agent performance, each agent keeps track of four different utility values: the *defection score* (DS) incurred by an agent who defects, the *punishment score* (PS) incurred by an agent who punishes or metapunishes another (as a result of an enforcement cost, and the *no punishment score* (NPS) incurred by an agent who does not punish another when it should, and is consequently metapunished. In addition these are combined into a total score (TS).

In this context, we can consider the algorithms used in our simulation, in two phases, as represented in Algorithms 2 and 3, called by Algorithm 1. More precisely, in Algorithm 2, each agent has various defection opportunities (o), and defects if its boldness is greater than the probability of its defection being seen. if an agent defects (Line 3), its DS increases by a *temptation payoff*, T

Algorithm 2 $\text{interact}(T, H, P, E)$

```

1. for each agent  $i$  do
2.   for each opportunity to defect  $o$  do
3.     if  $B_i > S_o$  then
4.        $DS_i = DS_i + T$ 
5.       for each agent  $j : j \neq i$  do
6.          $TS_j = TS_j + H$ 
7.         if  $\text{see}(j, i, S_o)$  then
8.           if  $\text{punish}(j, i, V_j)$  then
9.              $DS_i = DS_i + P$ 
10.             $PS_j = PS_j + E$ 
11.          else
12.            for each agent  $k : k \neq i \wedge k \neq j$  do
13.              if  $\text{see}(k, j, S_o)$  then
14.                if  $\text{punish}(k, j, V_j)$  then
15.                   $PS_k = PS_k + E$ 
16.                   $NPS_j = NPS_j + P$ 

```

(Line 4), but it *hurts* all others in the population, whose scores decrease by H (line 6), where H is a negative number that is thus added to the score. If an agent cooperates, no scores change. DS thus determines whether an agent should increase or decrease boldness in relation to its utility.

However, each hurt agent can in turn observe the defection and react to it with punishment that is probabilistic to its vengefulness. Punishment and metapunishment both have two-sided consequences: if an agent j sees agent i defect in one of its opportunities (o) to do so, with probability S_o (Line 7), and decides to punish it (which it does with probability V_j ; Line 8), i incurs a punishment cost, P , to its DS (Line 9), while the punishing agent incurs an enforcement cost, E , to its PS (Line 10). Note that both P and E are negative values, so they are added to the total when determining an overall value. If j does not punish i , and another agent k sees this in the same way as previously (Line 13), and decides to metapunish (Line 14), then k incurs an enforcement cost, E , to its PS , and j incurs a punishment cost P to its NPS .

In the learning phase, in Algorithm 3, and as mentioned above, each agent uses the various scores to determine how to improve its actions in the future. At the beginning of the learning procedure, the agent calculates its total score by combining all the other scores. In order to ensure a degree of exploration (similar to mutation in the original model's evolutionary approach, to provide comparability), we adopt an *exploration rate*, γ , which regulates adoption of random strategies from the available strategies universe (Line 8).

If the agent does not explore, then if defection is the cause of a low score (Line 12), an agent decreases its boldness, and increases it otherwise. Similarly, agents increase their vengefulness if they find that the effect of not punishing is worse than the effect of punishing (Line 22), and decrease vengefulness if the situation is reversed. As both PS and NPS represent the result of two mutually

Algorithm 3 learn(γ, δ)

```

1.  $Temp = 0$ 
2. for each agent  $i$  do
3.    $TS_i = TS_i + DS_i + PS_i + NPS_i$ 
4.    $Temp = Temp + TS_i$ 
5.  $AvgS = Temp/no\_agents$ 
6. for each agent  $i$  do
7.   if  $TS_i < AvgS$  then
8.     if explore( $\gamma$ ) then
9.        $B_i = random()$ 
10.       $V_i = random()$ 
11.     else
12.       if  $DS_i < 0$  then
13.         if  $B_i - \delta < 0$  then
14.            $B_i = 0$ 
15.         else
16.            $B_i = B_i - \delta$ 
17.       else
18.         if  $B_i + \delta > 1$  then
19.            $B_i = 1$ 
20.         else
21.            $B_i = B_i + \delta$ 
22.       if  $PS_i < NPS_i$  then
23.         if  $V_i - \delta < 0$  then
24.            $V_i = 0$ 
25.         else
26.            $V_i = V_i - \delta$ 
27.       else
28.         if  $V_i + \delta > 1$  then
29.            $V_i = 1$ 
30.         else
31.            $V_i = V_i + \delta$ 

```

exclusive actions, their difference for a particular agent determines the change to be applied to vengefulness. For example, if $PS > NPS$, then punishment has some value, and vengefulness should be increased. As indicated previously, this is covered in more detail in [8], but we will provide no further details here.

3 Imposing Topologies on Metanorms

Axelrod's model is interesting and valuable in examining how norms can be established in a population of agents. Using our simulation model, we are able to match Axelrod's results (and in fact improve on them, since Axelrod's model fails for extended runs of the simulation, as demonstrated by [7]). In a fully connected network (in which each agent is connected to every other agent), matching Axelrod's initial configuration, we get the results shown in Figure 1,

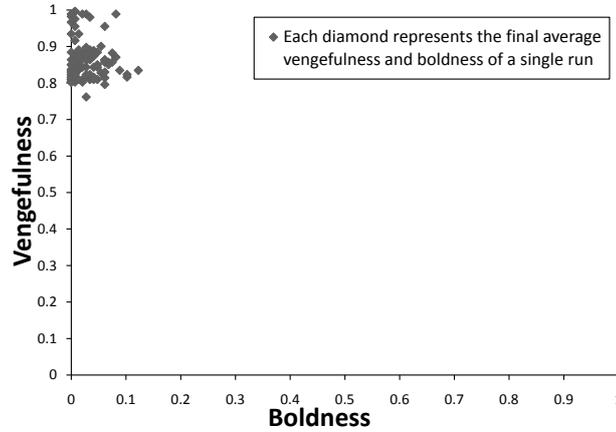


Fig. 1. Strategy improvement

and detailed in [8]. This provides a valuable illustration of the value of norms and the use of metanorms to avoid norm collapse in a system in which there is no central control, but Axelrod’s model omits consideration of some important aspects. In particular, in real-world computational domains, such as peer-to-peer and wireless sensor networks, the network of agents is not fully connected, with agents tending to interact with a small subset of others on a regular basis, yet it is only through such interactions that defection can be observed and punishment administered. Note that we restrict ourselves here to computational abstractions that apply to such environments rather than to physical or human networks.

Thus, while Axelrod’s model assumes a fully connected network, an unlikely and unreasonable assumption, other network topologies must instead be considered, reflecting different potential configurations of agents, in which agents are connected only to a subset of other agents, their *neighbours*. This constraint on connectivity between agents implies some adjustments to Axelrod’s model, as follows.

First, in Axelrod’s model it is assumed that an agent’s defection penalises all other agents in the population. The introduction of a topology enables us to restrict the penalty to only those agents with which the defector interacts. Second, in Axelrod’s model, agents are assumed to be able to observe the entire population. By introducing a topology, we employ a more realistic model in which an agent can only observe those agents with which it interacts. Third, punishment requires observation of misbehaviour. In Axelrod’s model, this requirement is implicit as it makes no meaningful distinction. However, by introducing constraints on observation and rendering the model more realistic, a further refinement is required: an agent can only punish a defector if the agent can observe the defector. In addition, an agent can only metapunish an agent that fails to punish a defector if it can observe both the defector *and* the agent that fails to punish

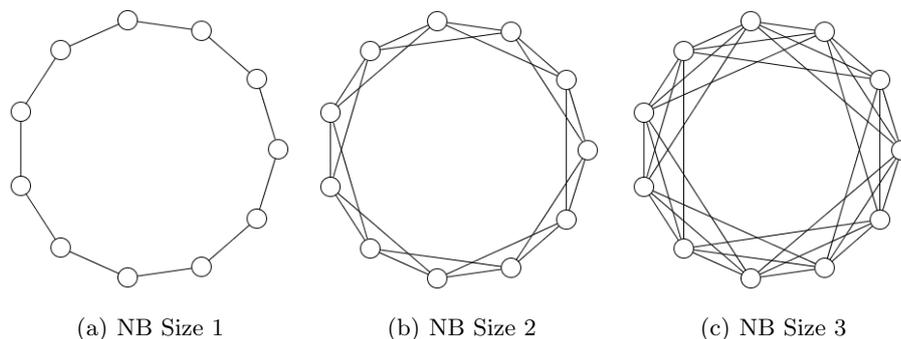


Fig. 2. Examples of lattice topologies

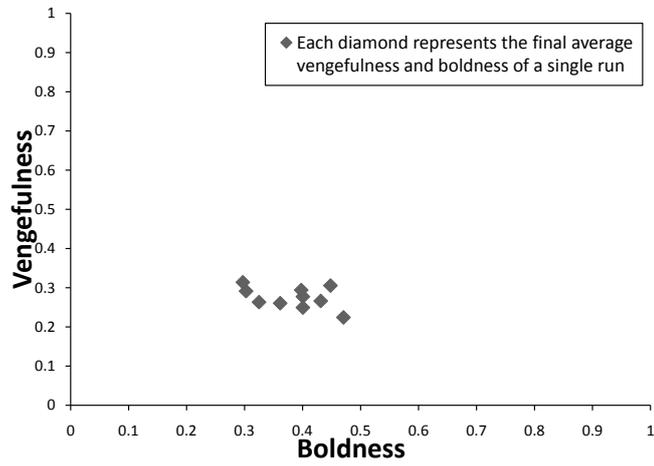
the defector. Finally, in order to enhance an agent's individual performance, it compares itself to others in the population before deciding whether to modify strategy. However, since agents can only observe their neighbours, these are the only agents they are able to learn from.

In consequence, the algorithms presented above are no longer adequate, and need to be changed as follows. First, in Algorithm 2, Line 5 needs to consider only agent i 's neighbours rather than all of the agents in the population, and Line 12 needs to consider only agent j 's neighbours. Then in Algorithm 3, the average score in Line 3, $AvgS$ should instead refer to the average score of the neighbourhood (that is, those agents to which agent i is connected. In this way, and with these simple modifications, our algorithms now address the needs of different topological structures.

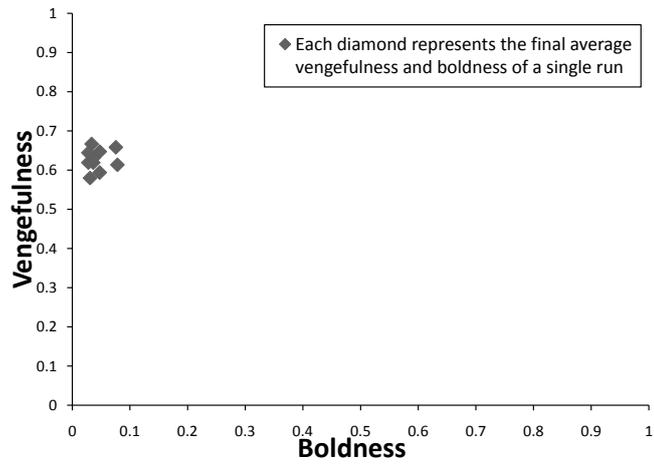
In what follows, we consider these modifications to the basic model in the context of different kinds of topologies, in particular small world models and scale-free networks. However, to start, we introduce lattices, since they provide the foundation on which small-worlds are based.

4 Metanorms in Lattices

A lattice (typically a simple ring structure) is perhaps the simplest network topology we consider, in particular, because it is also used as a base for more interesting and valuable topologies. In a (one-dimensional) lattice with neighbourhood size n , agents are situated on a ring, with each agent connected to its neighbours n or fewer hops (lattice spacings) away, so that each agent is connect to exactly $2n$ other agents. Thus, in a lattice topology with $n = 1$, each agent has two neighbours and the network forms a ring as shown in Figure 2(a). In a lattice topology with $n = 3$, each agent is connected to 6 neighbours, as shown in Figure 2(c).



(a) Lattice with neighbourhood size 1



(b) Lattice with neighbourhood size 3

Fig. 3. Smaller neighbourhoods in lattices

4.1 Neighbourhood Size

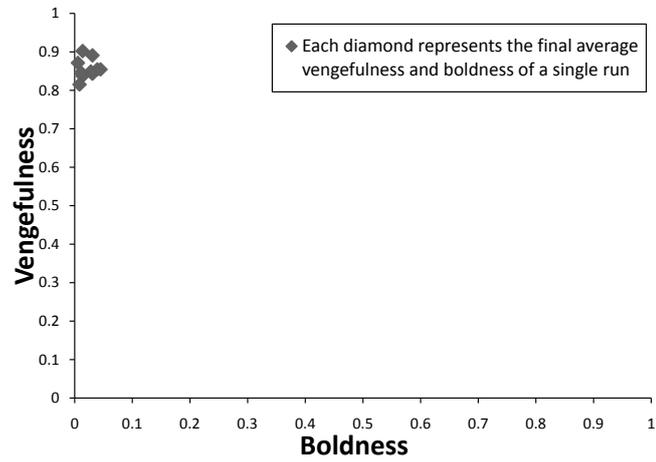
It is clear that, depending on the neighbourhood size, lattices may be more or less connected. Those with larger neighbourhood sizes are more similar to Axelrod’s fully connected model; our hypothesis is that as the neighbourhood size increases, the greater connections between agents enable punishment and metapunishment to become more effective in reducing boldness and increasing vengefulness. In order to investigate this hypothesis, we ran several experiments.

In our first set of experiments, we used 51 agents (so we have an even number, plus one, to account for the $2n$ neighbours plus our original agent), and varied the neighbourhood size between the least connected lattice (the ring topology) and the most connected lattice ($n = 25$). Each experiment involved 10 separate runs, with each run comprising 1,000 timesteps. for a particular neighbourhood size.

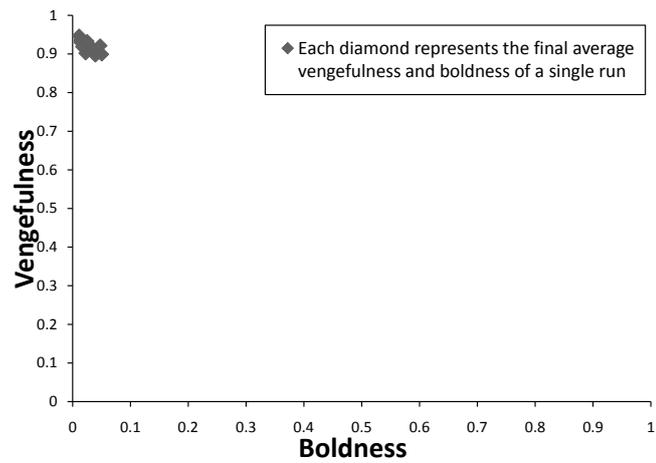
For the least connected lattice (n of 1), no norm is established, as runs ended in both relatively low boldness and relatively low vengefulness (see Figure 3(a)). In this case, though agents rarely defect, they also rarely punish a defection. This constitutes an unstable situation in which defecting could be a rewarding behaviour for agents as it is relatively unlikely to be penalised. However, increasing the neighbourhood size slightly to 3 (Figure 3(b)) has a noticeable impact on the results, as the boldness of the population drops almost to 0, which means that agents do not defect. While the level of vengefulness increases, it is still not at a level that can be considered to correspond to norm emergence, since agents might still not punish a defection without being metapunished for not doing so.

In addition, increasing the neighbourhood size to 13 has the same effect on boldness and a stronger effect on vengefulness (see Figure 4(a)), as vengefulness increases further, and almost to its maximum, of 1, when the neighbourhood size of 19 is used (see Figure 4(b)). These results suggest that increasing neighbourhood size strengthens norm emergence, by virtue of agents being more willing to punish norm violators. In seeking to provide more detail for analysis, the results of all runs were averaged, and shown on the graph in Figure 5, with neighbourhood size plotted against boldness and vengefulness. This shows that a neighbourhood size as small as 2 is enough to maintain boldness near 0, indicating that agents do not defect except when they *explore* as a result of sometimes adopting random strategies (introduced for comparability with Axelrod’s model). Conversely, increasing the neighbourhood size has a major impact on vengefulness, until the neighbourhood size reaches around 15 (at which point an agent is connected to half the population) when it brings only very minor change. This is because, in a poorly connected environment, agents that do not punish defection can more easily escape metapunishment than in a more connected environment.

As we hypothesised, increasing neighbourhood size brings a corresponding effect on the strategy of agents (in terms of boldness and vengefulness). Only the most poorly connected lattices have moderate levels of boldness, with vengefulness increasing monotonically over a longer period before it stabilises at a level consistent with norm establishment. The connections between agents give rise



(a) Lattice with neighbourhood size 13



(b) Lattice with neighbourhood size 19

Fig. 4. Larger neighbourhoods in lattices

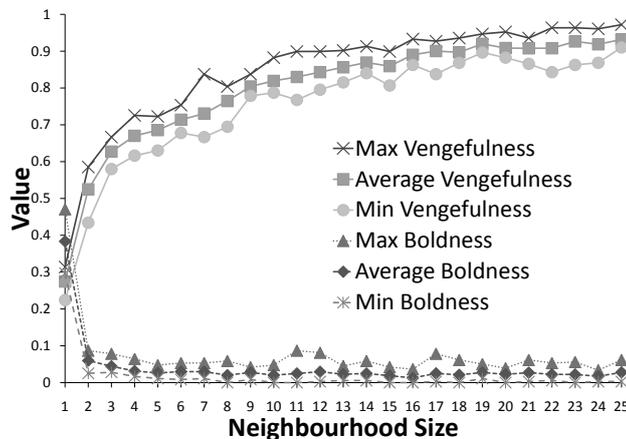


Fig. 5. Lattice: impact of neighbourhood size on final B and V

to this behaviour, with an increase in connections providing more opportunities for agents to respond to defectors appropriately.

4.2 Population Size

Now, if we increase the population size while keeping the neighbourhood size static, we decrease the relative number of connections among the overall population. This suggests that convergence to norm establishment should decrease, in line with the results obtained above. In the second set of experiments, therefore, the neighbourhood size was fixed and the population size varied between 51 and 1,001 agents. However, the results obtained, shown in Figure 6 for a neighbourhood size of 3 (though other values gave similar results), are not as expected, and suggest that increasing the population size has no effect on the rate of norm emergence, as all runs for all sizes of population end almost with the same level of boldness and vengefulness.

These results suggest that norm emergence in a community of agents that interact in a lattice is not affected by total population size but by neighbourhood size. By increasing the number of neighbours, norm establishment becomes more likely, irrespective of the size of the population. In other words, the likelihood of norm establishment is governed by the total amount of punishment that could potentially be brought upon a defector or an agent failing to punish a defector, which may be termed the *potential peer pressure* of a lattice. This is because such lattices essentially comprise multiple overlapping localities in which agents are highly connected: via punishments, the agents in these localities impose a strong influence on their neighbours. Increasing the population size simply increases the number of such overlapping regions.

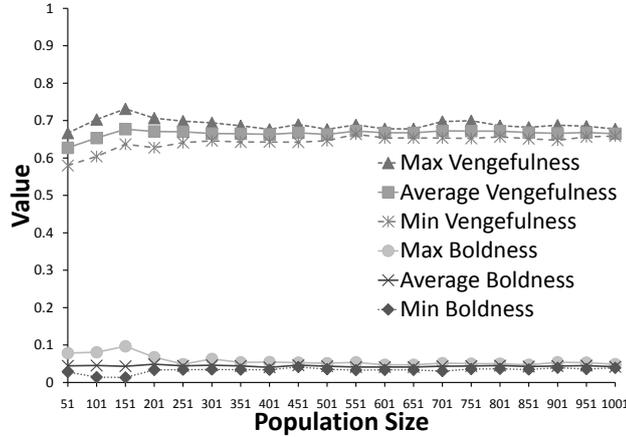


Fig. 6. Lattice: impact of population size on final B and V (neighbourhood size = 3)

5 Metanorms in Small Worlds

While lattices are regular structures, as opposed to random structures, Watts and Strogatz noted that many biological, technological and social networks lie somewhere between the two: neither completely regular nor completely random [16]. They instead proposed *small world networks* as a variation of lattices in which agents are connected to others n or fewer hops (on the ring) away, but with some of the connections replaced by connections to other randomly selected nodes in the network, in line with some specific rewiring probability (RP).

Thus, while lattices essentially create overlapping localities of well connected agents (since agents are connected to $2n$ agents immediately surrounding them), the effect of small worlds is to break these connections. Though the number of connections does not change, the locality effect does, since there may no longer be localities of well connected agents, but instead agents with some connections to their local neighbours, and some connections to others elsewhere in the network. As these local regions break down, the strong influence of an agent's local neighbours, causing compliance with norms, should also break down because of the more sparse connections.

To verify this hypothesis, we investigated the impact of the rewiring probability by running the model with different values, in populations of 51 agents, for different neighbourhood sizes of 3 and 5. The results of the experiment with a neighbourhood size of 3 are shown in Figure 8, which indicates that increasing the RP decreases the final average vengefulness in the population. With a neighbourhood size of 5 the results are similar (not shown).

This is because, as a result of rewiring, agents no longer affect just their locality, but now affect agents that are much further away, consequently requiring establishment of the norm in multiple localities. For example, in the case of neighbourhood size of 3, it is clear that not only is the norm not established,

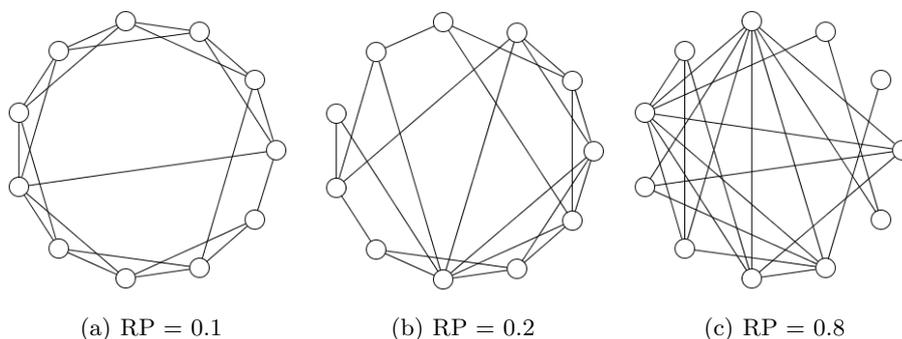


Fig. 7. Examples of small worlds

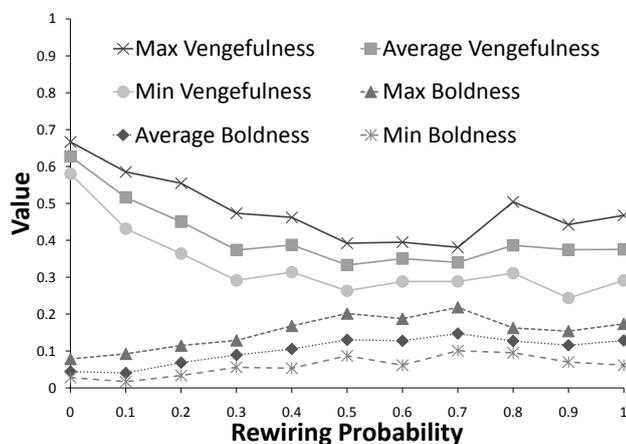


Fig. 8. Small world: impact of rewiring on final B and V (where neighbourhood size, $n = 3$)

but as the RP rises above small values, the trend moves further away from establishment, since the connections of agents are increasingly rewired, giving a locality effect similar to lattices with a neighbourhood size of 2 (discussed in Section 4.1). In addition, rewiring to other agents further away brings the need to establish the norm in all those localities to which an agent is connected, making it much more difficult.

In term of boldness, it is clear from the results that the RP of small worlds does not impact on the level of defection in the population since, independently, boldness remains very low, indicating that agents are very unlikely to defect.

5.1 Neighbourhood Size and Rewiring

As discussed in Section 4.1, increasing neighbourhood size causes an increase in vengefulness in lattices. In seeking to understand the impact in small worlds, we repeated the lattice experiments in this new context, for different values of the RP. Results for a rewiring probability of 0.4 are shown in Figure 9 (with results for other values of the RP being similar in trend), again showing that neighbourhood size increases vengefulness. However, note that, in comparison to lattices, vengefulness in small worlds is lower for the same neighbourhood size. This is because the agents must now respond to defections in different regions of the network, where there is less influence on behaviour, and thus potentially incurring greater enforcement costs.

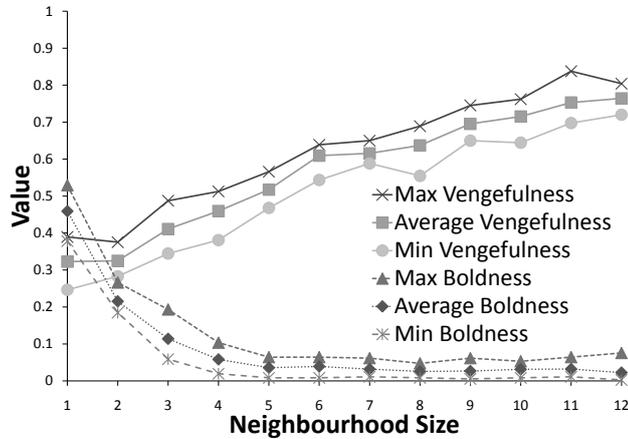


Fig. 9. Small world: impact of neighbourhood size on final B and V (RP=0.4)

5.2 Population Size and Rewiring Probabilities

Population has been shown to have no effect on norm establishment in lattices due to the *potential peer pressure* arising from the multiple overlapping localities. However, since these concentrated local regions of connected agents are weakened in small worlds, we repeated the previous experiments to determine the effect with RPs of 0.2, 0.4, 0.6, 0.8 and 1.0, and n of 5. The results indicate that boldness is not affected by the changes of the population size as it is always close to zero (not shown), but vengefulness decreases as the RP increases. More specifically, when the RP is 0.2, increasing the population size has little effect, as shown in Figure 10. However, for the other RP values, increasing the population size decreases vengefulness. Again, this is due to rewiring breaking down the strong locality effect, and this is magnified with increasing population sizes,

since there is a greater opportunity for connections to other localities, causing a greater cost for agents seeking to bring about norm establishment in all these localities at once.

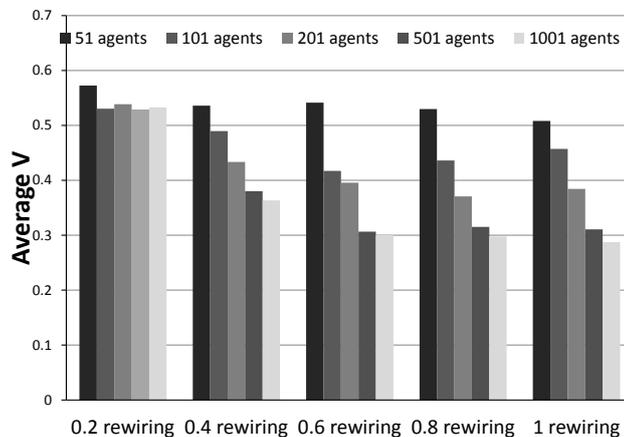


Fig. 10. Small worlds: impact of rewiring and population size on final vengeance (where neighbourhood size $n = 5$)

6 Related Work

There has been much work that focusses on the issue of norm emergence in societies of interacting agents. However, most of this concentrates on analysing norm emergence over fully-connected networks [2, 12, 13, 15], and it was only relatively recently that attention shifted towards the effect of the structure of these societies. In this section, we review that part of the literature that does address these concerns.

In a particular effort, Delgado et al. [4] study the emergence of coordination in scale-free networks. Their study involves an interaction model of a multi-agent system, by which they try to analyse how fast coordination can spread among agents. Coordination here is represented through agents being in the same state, which is achieved when 90% of the agents do so. The framework they use is rather simple, however: an agent makes a choice between two different actions and they receive a positive payoff if they both choose the same action, or a negative payoff if their actions are different. Agents record the outcome of taking each of the two actions and pick the action with the better outcome for next interaction. The results of the work demonstrate that coordination can indeed be achieved over scale-free networks, but in a rather restricted setting.

Similarly, Sen et al. [11] use a game to investigate norm emergence over lattices and scale-free networks. In particular, they analyse the effect of increasing

the number of actions available to agents, as well as the effect, on the speed of norm emergence, of increasing the number of agents in both scale-free networks and lattices. Their results suggest that both increasing the number of actions *and* increasing the number of agents causes a delay to the norm emergence in the population over a scale-free network. Similarly, norm emergence in lattices is much slower when agents have a larger set of actions to choose from, or when the number of agents in the population is increased. Overall, their analysis shows that, for a small set of actions, it is faster for a norm to spread in a ring than in other topologies, followed by fully connected structures, and then scale-free networks. In contrast, for a large set of actions, it turns out that this is much faster in scale-free networks than in rings and fully connected structures.

As we have suggested, the models used in these previous pieces of work are relatively unsophisticated, with only two agents involved in an interaction, and reward values remaining fixed and not changing during the game. In response, Villatoro et al. [14] adopted the same concept of two-agent interactions, but introduced the notion of the reward of an action being determined through the use of the memory of agents, thus adding some dynamism to the model. Here, the reward of a certain action is determined by whether the action represents the majority action in both agents' memories, and the reward is proportional to the number of occurrences of this majority action in their memories. However, it is not clear from where these rewards derive nor who applies them, as agents only have access to their memory. With regard to interaction networks, their work illustrates that increasing the neighbourhood size of a lattice accelerates norm emergence. In contrast, in the case of scale-free networks, norms do not emerge using the basic model. This is because of the development of *sub-conventions* that are persistent and hard to break, and which prevent the whole population from converging towards a single convention. A solution to this problem was found by giving *hub* agents (those with the majority of connections to others) more influence on the reward function.

Savarimuthu et al. [9] analyse the effect of advice on norm emergence over random and scale-free networks. For this reason, they use the *ultimatum game* in which two agents must decide how to share a certain amount of money. One agent offers a particular division of the money to the other and, if the second agent agrees, then the money is divided between the two agents according to this proposal. If the second agent does not agree, both agents gain nothing. Here, each agent has a personal norm that defines its proposal strategy and, in addition, agents are able to request advice about their proposal strategy from a *leader* agent that is believed to have the best performance in the neighbourhood. However, agents are capable of accepting or refusing the advice according to their autonomy level. The results obtained in this work show that norm emergence increases in speed over both random and scale-free networks with an increase in the average degree of connectivity.

Our work is rather different to these previous efforts in that we have investigated a more sophisticated model. In addition, we are not restricted to only two agents, and consider arbitrary numbers of them, since any agent's actions can be

observed by all of its neighbours. These neighbours can in turn react by choosing to punish or to avoid doing so, potentially generating further metapunishments by other observing agents. Finally, sanctions applied in our case are dependent on the decisions of all agents that observe a violation, thus making them change with the number of agents involved.

7 Conclusions

In this paper, we have investigated mechanisms that encourage norms to emerge in communities of self-interested agents, without interference of a central or outside authority, under the realistic constraint that agents can only influence one another if they regularly interact. Based on Axelrod's seminal work, our model's substantial novel extension examines the impact of different types of topologies of interaction on norm emergence. Our results show that in circumstances in which *each* agent regularly interacts with a small number of other agents, as in lattices and small worlds, Axelrod's mechanisms to encourage norm emergence remain largely effective. More precisely, it is very effective for lattices, but its effectiveness varies with the rewiring probability in small worlds. Moreover, we have demonstrated that, given fixed penalties, for lattices, the effectiveness of Axelrod's approach only depends on the number of neighbours of each agent, *not* on the total population size. For small worlds, increasing the population size with a high rewiring probability decreases vengefulness, constraining norm emergence significantly. Thus, topology must be considered: in the case of a lattice or a small world, Axelrod's proposed approach will be effective for sufficiently large neighbourhood sizes.

References

1. R. Axelrod. An evolutionary approach to norms. *American Political Science Review*, 80(4):1095–1111, 1986.
2. M. Boman. Norms in artificial decision making. *Artificial Intelligence and Law*, 7(1):17–35, 1999.
3. A. P. de Pinninck, C. Sierra, and W. M. Schorlemmer. Friends no more: norm enforcement in multiagent systems. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pages 640–642, 2007.
4. J. Delgado, J. M. Pujol, and R. Sangüesa. Emergence of coordination in scale-free networks. *Web Intelligence and Agent Systems*, 1:131–138, 2003.
5. J. M. Epstein. Learning to be thoughtless: Social norms and individual computation. *Computational Economics*, 18(1):9–24, 2001.
6. F. Flentge, D. Polani, and T. Uthmann. Modelling the emergence of possession norms using memes. *Journal of Artificial Societies and Social Simulation*, 4(4), 2001.
7. S. Mahmoud, N. Griffiths, J. Keppens, and M. Luck. An analysis of norm emergence in axelrod's model. In *NorMAS'10: Proceedings of the Fifth International Workshop on Normative Multi-Agent Systems*. AISB, 2010.

8. S. Mahmoud, N. Griffiths, J. Keppens, and M. Luck. Overcoming omniscience in axelrod's model. In *Proceedings of the IAT Workshop on Coordination, Organizations, Institutions and Norms*, 2011.
9. B. T. R. Savarimuthu, Stephen Cranefield, Maryam Purvis, and Martin Purvis. Role model based mechanism for norm emergence in artificial agent societies. In *Coordination, Organizations, Institutions, and Norms in Agent Systems III, COIN 2007 International Workshops*, volume 4870 of *Lecture Notes in Computer Science*, pages 203–217. Springer, 2008.
10. B. T. R. Savarimuthu, M. Purvis, M. Purvis, and S. Cranefield. Social norm emergence in virtual agent societies. In M. Baldoni, T. C. Son, M. B. van Riemsdijk, and M. Winikoff, editors, *Declarative Agent Languages and Technologies VI*, volume 5397 of *Lecture Notes in Computer Science*, pages 18–28. Springer, 2009.
11. Onkur Sen and Sandip Sen. Effects of social network topology and options on norm emergence. In *Proceedings of the Fifth International Conference on Coordination, Organizations, Institutions, and Norms in Agent Systems*, pages 211–222, 2010.
12. S. Sen and S. Airiau. Emergence of norms through social learning. In *IJCAI 2007: Proceedings of the Twentieth International Joint Conference on Artificial Intelligence*, pages 1507–1512. Morgan Kaufmann Publishers Inc., 2007.
13. Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: off-line design. *Artificial Intelligence*, 73(1-2):231–252, 1995.
14. D. Villatoro, S. Sen, and J. Sabater-Mir. Topology and memory effect on convention emergence. In *Proceedings of the 2009 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technologies*, pages 233–240. IEEE, 2009.
15. A. Walker and M. Wooldridge. Understanding the Emergence of Conventions in Multi-Agent Systems. In V. Lesser, editor, *Proceedings of the First International Conference on Multi-Agent Systems*, pages 384–389. MIT Press, 1995.
16. D. J. Watts and S. H. svaratz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.
17. T. Yamashita, K. Izumi, and K. Kurumatani. An investigation into the use of group dynamics for solving social dilemmas. In P. Davidsson, B. Logan, and K. Takadama, editors, *Multi-Agent and Multi-Agent-Based Simulation*, volume 3415 of *Lecture Notes in Artificial Intelligence*, pages 185–194. Springer, 2005.