

Cooperation in Growing Communities

Rowan Martin-Hughes

Abstract As communities grow in size over time from just a few people to hundreds and then thousands, members frequently find that they feel less involved, that the community lacks relevance, and that their trust in the community as a friendly place is gone. A prime example of this is online message boards or other communities developed around social interaction which are renowned for becoming bogged down in endless arguments and spamming as they increase in size. The same ideas apply to online trading systems such as eBay which require a far higher degree of trust and reliability. We follow a game theoretic model of frequent interactions over time between reactive agents to examine the conditions under which a population is likely to find a set of strategies which allow them to cooperate a sufficient percentage of the time to remain viable.

1 Introduction

Whether through a polite nod when passing on the street or through a more tangible exchange of items, people interact in many ways which have the potential to be mutually beneficial. These are transactions which have no inherent obligation for reciprocity but for which society has evolved strict laws, subtle peer pressure, or other mechanisms to encourage appropriate behaviour for the benefit of all. A society in which people feel confident that these potentially positive exchanges *will* take place is then by definition one in which there is a high degree of trust.

As the number of online and other long distance interactions increases we face many new challenges in maintaining a level of trust. A game theoretic consideration of this social shift would suggest to us that without repeated interaction the ability for one person to encourage pro-social behaviour in another through potential punishment is vastly diminished, and this reduction in consequences for anti-social

Rowan Martin-Hughes
The Australian National University, e-mail: rowan.martin-hughes@anu.edu.au

behaviour may lead to undesirable outcomes. In thinking about online interactions specifically we can see a number of additional challenges as anonymity via cheap pseudonyms [8] is relatively easy to achieve and external judicial measures are often less readily available. The challenges are particularly clear when users are interacting at a distance, with the increased likelihood of added difficulty as differing cultural backgrounds make the probability of misunderstandings higher.

The question then is how critical each of these factors are to the emergence of trust from a community of users? We will approach this by simulating a group of agents and studying the conditions under which trust is able to emerge, thus demonstrating some broadly qualitative results.

2 Background

We begin by discussing a study by Lomborg [11] which showed how a community of agents playing the iterated Prisoner's Dilemma game could evolve over time from consisting of all defectors to a state where the average payoff more closely approached the ideal situation where cooperation is the norm, without any outside interference.

Early competitions between strategies consistently saw the tit-for-tat strategy outperforming others [2], but once we include noise tit-for-tat easily falls into a vicious cycle of defection, and these kinds of competitions did not allow much scope for the development of different strategies over time and across the population as a whole.

Lomborg simulated reactive agents playing iterated Prisoner's Dilemma games and analysed the strategies being used as the population changed over time. Over long periods in which the average payoff remained stable at close to the payoff of mutual cooperation, the agents were seen to be playing a mix of strategies which conformed to a configuration made up of two broad types of strategy. These were described as a *nucleus* of forgiving cooperative types, and a *shield* of strategies who would play cooperatively if their opponent did but who would aggressively punish any defectors. The resistance of this mixture of strategies to defection was due to the nucleus choosing to cooperate a vast majority of the time even under noisy conditions, while the shield was just sufficiently numerous and punitive against defection that an invading strategy involving unprovoked defection would receive a lower payoff over time than a regular part of the population.¹

With such a population breakdown, Lomborg found that the greatest vulnerability (aside from an overwhelming influx of defectors at once) was through complacency. At times where no invading strategies of defection exist the nucleus tended to grow at the expense of the shield, which left inadequate defence against an inevitable eventual threat.

¹ This phenomenon of the nucleus and shield was also observed in our system over long runs where the system was stable, with many agents using a highly cooperative strategy followed by a wide variety of (mainly slightly more zealous) variants on tit-for-tat.

This pattern has similarities with various models proposed in the field of psychology to describe the different social value orientations that people subscribe to. A review by Au and Kwong [1] describes a number of studies which indicate different breakdowns of prosocial (altruistic and cooperative) choices compared to individualistic or competitive choices. There does appear to be a great deal of variation within such studies though, which could be explained by the difficulty of measuring these distributions. A human playing what we call a nucleus strategy might show that in different ways depending on the phrase of the question and what game they feel they are playing - if they feel that they are being tested on their ability to cooperate with the experimenter in appropriately punishing a hypothetical defector then their response might be different to their choice in an open system where they have been wronged by another person and have to decide whether to seek justice.

It has been suggested that human emotions play a key role in ensuring that the appropriate balance of strategy types is maintained [16], and further evidence for this angle of investigation comes from neurological studies which showed reward centres in many people's brains were activated by punishing defectors [10]. These concepts and the way in which the number of people who act upon different strategies correlates with the mixture of agent strategies in Lomborg's model suggest that further examining the result of agents playing these types of games could supply useful feedback on describing and enforcing human social norms.

2.1 Trust

Reactive agents such as these are too simple to have notions of trust; they simply pursue what they believe is the most effective individual strategy at each point in time. From an outside perspective though we can describe this type of community as being trustworthy if cooperation is chosen regularly in the iterated Prisoner's Dilemma game.

This will occur when a majority of agents believe that a highly cooperative strategy will lead to the highest expected payoff. Thus a good measure for this notion of trust is the simple utilitarian measure of the average utility each agents gains over time by playing these games. An average payoff for agents which is close to the payoff for mutual cooperation will hence represent a trusting community while an average payoff approaching the payoff for mutual defection will show a system in which trust is almost entirely lacking.

3 The Model and Parameter Discussion

We are examining a society of simple agents of quite bounded decision making who try to maximise their utility over the course of many interactions with other agents. Such agents can be defined as playing transaction games wherein each agent has

the possibility of spending some utility to give their opponent utility. In the generic form of such a game an agent has the choice of spending u_1 to give their opponent u_2 , and likewise their opponent has the choice of spending u_3 to give the agent u_4 utility as shown in Table 1.

Table 1 The Transaction Game Payoffs

Agent/Opponent	Cooperates	Defects
Cooperates	$-u_1 + u_4 / -u_3 + u_2$	$-u_1 / u_2$
Defects	$u_4 / -u_3$	0/0

The meaning of the utility transfer is deliberately left quite open. Most obviously it may refer to an actual trade of goods for money, but it could equally be an exchange of ideas on a message board where each party has the option of trying to show themselves as being intellectually superior to the other or to be polite.

In an trading system such as eBay the use of a bidding system or other such mechanism generalised as *market forces* means that it makes sense for both agents to be considering this trade legitimately. Likewise in communications over a message board it is likely to cost relatively little effort to be polite (cooperate) than what the other person gains by feeling respected. Thus we can reasonably assume $u_1 < u_2$ and $u_3 < u_4$ is the standard situation leaving us with a (potentially asymmetric) Prisoner's Dilemma game.

With this we have sufficiently strong parallels to the model of Lomborg to consider the circumstances in such a modified community under which cooperation emerges to a satisfactory level. The main immediate difference is that most frequently the agents will only play once with other agents rather than iteratively, which makes the job of gaining cooperation dramatically harder. However, if they remember the actions their opponents took previously then over time they are still slowly playing iterated Prisoner's Dilemmas against each of the other agents.

Although it does not conform precisely to the structure of a transaction game, we use the same common formulation of the Prisoner's Dilemma as Lomborg, with payoffs given in Table 2.

Table 2 Default Payoffs

Agent/Opponent	Cooperates	Defects
Cooperates	3/3	0/5
Defects	5/0	1/1

3.1 Agents

Agents must be able to learn, innovate, have sufficiently complex strategic choices so as to have meaning while still being simple enough to model, and the system should include the possibility of noise. We use the same initial model of these factors as Lomborg did before describing how each of these and a number of other salient factors can be modified in more detail.

As a basis for this we follow Lomborg in defining the agents as Moore machines [4] who base their actions on a standard initial choice followed by actions dependent on the history of their opponent against them.

The reactive decision for each agent in a game is then based only on the most recent n moves of their opponent against them, giving them $\sum_{a=0}^n 2^a = 2^{n+1} - 1$ different situations for which they have a chosen action. For example if $n = 3$ there are 15 such situations as their opponent may have played 0, 1, 2, or 3+ times against them previously and could have any combination of cooperation or defection in each of those cases. Each agent thus has $2^{2^{n+1}-1}$ possible strategies, or 2^{15} in this case. The order in which this represents their response to the actions of the opposing agent is shown in Table 3. This equates to quite bounded rationality, but it will be seen later that manipulation of the other variables in the simulation can equate to greater deviousness of agents who may wish to be malicious.

Table 3 The opposing actions (least recent to most recent) which each point in an agent’s strategy corresponds to. For readability, strategies will be broken down into segments corresponding to the number of moves their opponent has previously played. The classic tit-for-tat strategy is then represented by these types of agents as {C,DC,DCDC,DCDCDCDC}.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
initial	D	C	DD	DC	CD	CC	DDD	DDC	DCD	DCC	CDD	CDC	CCD	CCC

As an example, consider the strategy {C,DC,CCCC,DDCCDDCC}. This agent cooperates first, then does whatever their opponent did the first time they played. The third time they play the same opponent they cooperate regardless and on the fourth and subsequent game they mimic the action of their opponent from two games previously.

3.2 Initial Conditions and Timeframe

The initial conditions we use involve all agents being in a state of defection. This is not something we are largely forced to accept in a real system, but it provides a consistent platform to compare parameters as well as the greatest challenge for such social solutions as our agents can find.² It can take a lengthy time period to escape

² If the parameters of the system show cooperation will develop and we already have cooperation then we expect it to be maintained over time, with perhaps small lapses which are quickly corrected

from this state of mutual defection to a more cooperative situation, and the time-frame of the simulations reflects this in being on a scale of hundreds of thousands of interactions.

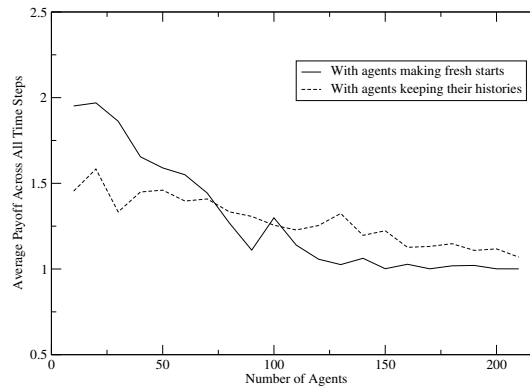
In the real world we can assume that humans have developed mechanisms such that their original strategies on joining a group often closely mirror the kinds of strategies we develop over time, either through learning from previous similar communities or via an evolutionary explanation. Furthermore, it seems reasonable to assume that people can make individually more intelligent decisions than our reactive agents and are able to reason through this process themselves, thus choosing a strategy they feel will be in equilibrium after all of the people involved have gone through such deliberation. Due to these factors, the results obtained have little quantitative meaning in terms of the time taken to reach a given community state for instance, but the long term qualitative effects of changing the basis of the community remains relevant.

3.3 Number of Agents

The number of agents in a large-scale open system is by definition virtually unlimited, and if each agent by default interacts randomly with any other then most agents will only interact with other agents only once or twice if at all. Under these circumstances our Prisoner's Dilemma scenario is not iterated between pairs of agents at all and there is no ability to punish defection, so things do not tend to progress beyond the all defection stage.

Figure 1 shows how cooperation between agents decays to the mutual defection payoff of 1 as the population size approaches even the small number of 200 users.³

Fig. 1 This figure shows the way that it becomes more difficult for cooperation to flourish for any length of time even in small populations. It also compares the effect of agents maintaining their histories as opposed to letting them be cleared as discussed at the end of Section 3.6.



for. Conversely, if we find that our parameters show that cooperation will be highly unstable at best and we start with cooperative agents we expect that it is likely to degenerate.

³ As a curious aside, we can see that cooperation is already scarce around 150 users which is Dunbar's Number [7], the number which anthropologist Robin Dunbar postulated as being where

Even at the very small population sizes we can see that the average payoff gained is a long way from the mutual cooperation payoff of 3. While these small populations are often in a cooperative state it requires little for them to temporarily fail and be dominated by defection again in an irregular cycle, as a sample run over a long time period with 25 agents shows in Figure 2.

3.4 Interaction Patterns

In many social interactions user activity is distributed according to a power law [13], and other aspects of network structure not covered by fully random interactions. Figure 3 shows some single sample runs of the average payoff in the system over time as we vary the frequency with which agents interact with the 10 agents to either side of them rather than with a random agent from the population at large.⁴ With localised interaction eventually two agents willing to try cooperating will share a neighbourhood allowing them to outperform defectors [15], whereas without this property there are never enough repeated interactions between cooperators and the cooperators quickly switch back to a defector strategy.

Because of the poor performance of any system with a large population without this kind of property or a reputation system of the form described later, we assume that agents act according to a locally connected world in order to judge the effect of other parameters.




Fig. 2 A sample run over a long period of time with 25 agents, showing how cooperation is frequently found but rarely stable for any length of time.

a group of people loses social cohesion essentially because no one person is able to know all of the others.

⁴ If we imagine all of the agents to be placed around a circle, then these agents show a regular social order in interacting with local “friends” some portion of the time and anyone else the rest of the time. When looking at actual people we may also be able to exploit the small-world properties [3] of infrequent long distance links to improve system efficacy.

3.5 History

Figure 4 shows sample runs where agents maintained a history size of 1, 3 or 5 past interactions with each other agent, and we can see how a short history translated into very quickly exploring the state space and finding a cooperative strategy, but also that it was completely unstable. With the limited strategies available to an agent with a history size of 1 this equates to essentially a cycle of dominant strategies across the agent population. Tit-for-tat fast becomes the dominant strategy, but is eroded over time by the equally good and more noise resistant cooperation strategy. This lasts until defection becomes the dominant strategic choice, quickly leading back to tit-for-tat replacing the now poorly performing cooperation strategies which results in defection performing poorly again as well. The dips that can be seen with the more complex strategies related to longer histories are examples of the same effect. Although it seems that a longer history size is likely to produce more consistent results, we used a history size of 3 for further tests due to the number of possible strategies with a large history size limiting our ability to closely observe strategic trends.

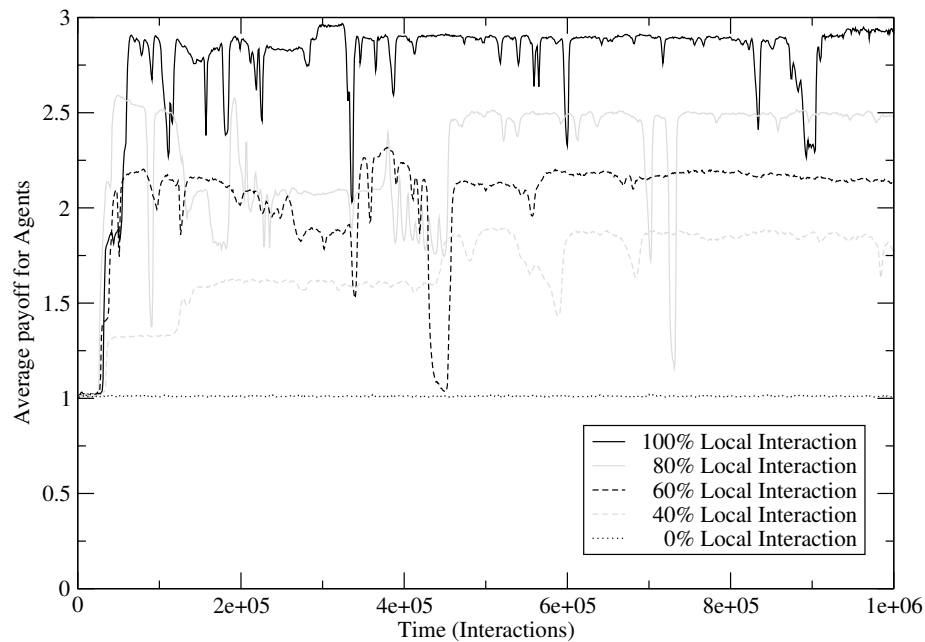


Fig. 3 A comparison of the average payoffs agents are able to receive over time in a population of 1000 agents over 100000 time steps, given varying levels of local vs random interaction.

3.6 Innovation and Learning

In order to explore the space of possible strategies, agents modify their strategies on an infrequent basis. They may firstly innovate and make either a small adjustment to their current strategy, less commonly jump to a completely random new strategy, or still more rarely they may revert to the primal strategy of always defecting.

More frequently than innovating, the agents will demonstrate learning through imitation [5] in which they compare their performance to another randomly selected agent. If they feel they are being outperformed by the other agent then they will change their strategy to match that of the other agent. This seems sub-optimal because they could seek out the very best performing strategy at each point in time or indeed perform some kind of search to determine the optimal strategy given the state the other agents are in and their beliefs about the future, but these considerations can reasonably be described as too costly for our agents (as they are for humans in large populations) to determine.

The numbers used have little significance from a qualitative point of view, but every 10th time step 5% of the agents consider imitating a different strategy and 0.1% of the agents innovate. Of the agents who change strategies to something new, 90% simply invert one part of their current strategy, 9% choose a completely random strategy, and the remaining 1% revert to pure defection.

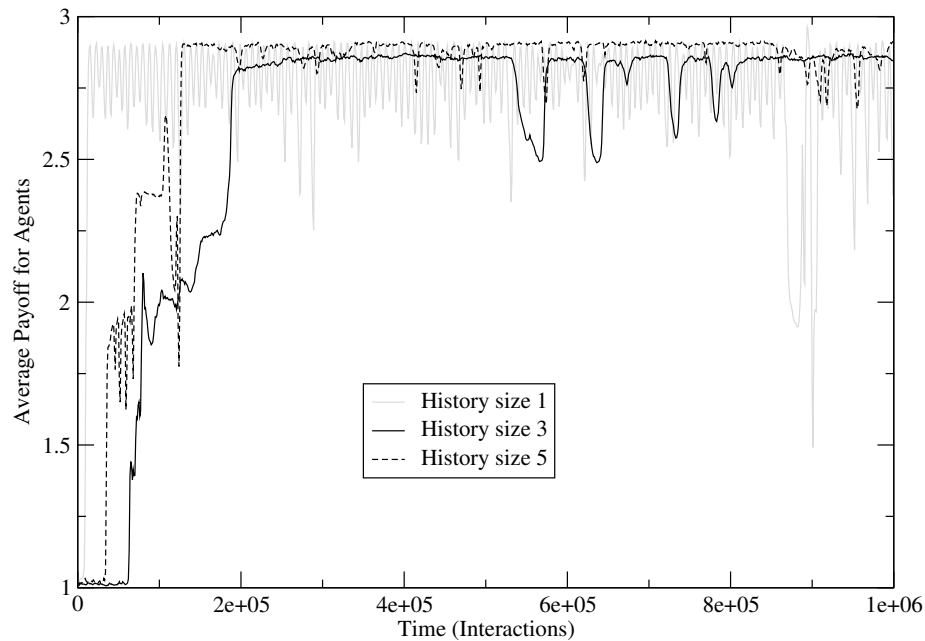


Fig. 4 This figure shows the differences between the amount of history used by agents when finding strategies for some single sample runs in a population of 1000 agents over a time period of 100000 interactions.

When agents decide to modify their strategy we find that we need to decide how they behave and are treated with respect to past actions. The first option is that they essentially start anew, forgetting the history of other players against them and being forgotten in turn as well as starting afresh in terms of their utility. This represents either a society willing to accept this, or the agent rejoining under a different name, or simply an agent leaving and being replaced. Alternatively, the agent may remember and be remembered, keeping their old utility and thus having their average payoff change slowly as they ease into the new strategy. This has the advantage of showing how an agent might modify their behaviour while still being observed, but it makes it difficult for agents to recognise whether the success of an agent belongs to their current strategy or a previous one, as well as making the initial part of the strategy less relevant over time.

Figure 1 shows how keeping the old history slightly dulls the ability of the population to find cooperative strategies in small populations, but also makes the presence of cooperation last for longer as the population grows. This seems to be as for a period after changing strategies to a more cooperative strategy even when it is unfavourable, other agents may emulate this if the previous strategy of defection was giving a higher payoff. In order to see the effects of other factors more clearly all agents clear their histories and start afresh, although in a real system there would be a mixture of these mechanisms.

3.7 Noise

We again follow Lomborg in considering the situation where the system is noisy but the agents do not know if their action has been misinterpreted. Hence when it occurs the payoffs stay the same as intended, but an agent has a small chance of putting the opposite action in their history.

Unsurprisingly noise makes the system less productive as in Figure 5, and based on this we globally assume a noise level of 0.01 although other combinations of conditions may allow us to tolerate higher levels of noise. A reduction in noise can occur in a trading system through encouraging better communication and more complete descriptions of items with photos. In a message board or comments system the goal of reducing communication difficulties is the same, so any structures which are able to discourage inflammatory comments on a general basis raising the tolerance for genuinely different opinions will have this kind of effect for instance.

3.8 Strategic Choice and the Folk Theorem

One of the randomly chosen strategies that can often be seen heading the list of worst performers early in a simulation run is {D,DD,DDDD,CDDDDDDDD} which is comically awful in the way it angers everyone by defecting repeatedly but then

curls up and plays dead under repeated punishment. While this seems a foolish choice under that state of the world, when playing against very forgiving opponents the rare apologetic confession after the opponent is sure they have to retaliate allows it to get away with more defections and a higher payoff in the long run.

In fact, by the Folk Theorem [9] every strategy has a counter strategy able to outperform it and thus invade the population over time. Without going into detail on this, the gist is that with the existence of credible threats as there is in the Prisoner's Dilemma *any* combination of strategies between fully rational players can be in equilibrium. Essentially this works because one player can always threaten the other with constant defection if they do not agree to play that strategy, which the other player can respond to in turn.

This vitally limits the possibility of stating some perfect population distribution which results in complete trustworthiness as any such society will be lacking a shield. Conversely, any society with an appropriate shield against defectors will see a constant dip in performance as agents in the shield will develop feuds over time due to noise. The important focus must remain on maintaining the conditions which allow trust to robustly return when defectors are able to invade in numbers for a period of time.

3.9 Frequency of Innovation and Imitation

Questions may arise as to why a successful agent changes strategy, or why a really unsuccessful agent may just keep on trying with a poorly performing strategy. Both of these see parallels in the bizarre behaviours that humans can exhibit at times, and they can also be seen as allegorical at times - a successful agent may simply leave the system and be replaced at the same time by some new strategy.

Of course, in a real trading system the payoffs for each individual transaction are unique, but for these agents it makes no practical difference over thousands of iterations to have individual payoffs instead of receiving an average every time as




Fig. 5 This figure shows the debilitating effect of increased noise on a community of 1000 agents over a period of 100000 time steps.

we allow *market forces* or *social norms* to ensure that the payoffs are individually sensible. In the trading domain we have what seems to be another difficulty arising from a large population size. When many agents bid on a hypothetical auction the marginal benefit of cooperation to the buyer is likely to decrease, making them less likely to cooperate. However, it also *increases* the marginal benefit of cooperation to the seller, and as many trading systems are set up such that only the seller really has the opportunity to defect in a significant way,⁵ this actually helps the chances of cooperation. In a truly competitive market where both buyers and sellers are competing, margins are likely to end up small and defection will thus in theory be more tempting.

The effect on the chances of cooperation as the relative payoffs change in this manner or some other manner has been studied for a long time beginning with Rapoport's index of cooperation $K = \frac{CvC - DvD}{DvC - CvD}$ based on the relative payoffs available to a player for the game with payoffs $\{CvD, DvD, CvC, DvC\}$ [14]. While this is not specifically useful aside from as a general guide, it is clear that an agent who receives payoffs of $\{0, 1, 300, 301\}$ instead of $\{0, 1, 3, 5\}$ will find it far easier to stay in a state of cooperation once they reach it. Even if such an agent only rarely plays an opponent who obliges them with cooperation, a frequent defector is unlikely to exist with a higher average payoff.

Unfortunately we also have to consider the problem of agents who are not honest traders, in that they agreed to the trading contract with no reason to carry it through in the first place. Whether by accident of timing or through intention to deceive from the beginning, their payoffs are likely to make defecting strictly optimal because the payoff for mutual cooperation will be lower than for mutual defection. We examine the effect of having a number of agents like this in the system in Section 4.

The complications that arise from either of these modifications are hard to gauge exactly, because in both cases the payoffs cease to be easily comparable between agents when they consider changing their strategy. More importantly from the perspective of evaluating the success of the system, it is hard to say what a good result is - some agents may simply have a higher capacity to gain utility and thus could be judged as more important, or equally we could compare the percentage of cooperation with no thought to the individual gains. When we add a second type of agents their payoffs are comparable so this philosophical question can be partially avoided, but utility comparisons between agents have long been a pitfall of game theoretic solutions to social problems.

4 Reputation Systems

The effectiveness of many systems is based around the ability to choose opponents and thus exclude players who defect, as the threat of being labelled as a defector and excluded is enough to make most rational agents cooperate. For instance, eBay's

⁵ By virtue of up front payment from the buyer rather than simultaneous transaction.

reputation system is effective on the basis of allowing other people to avoid a player with a poor reputation rather than as a directly punitive action. Research on certain types of low cost items indicates that reputation can have a negligible effect on the sale price but a greater effect on the probability of a sale although other studies on different types of items found mixed results [6]. The moderation system used by Slashdot essentially does this also as poorly ranked post(er)s are simply unseen by most users and are thus not able to play the game of exchanging opinions.

An equivalent point of view is that essentially a reputation system forces each user to play a second parallel iterated game against *the system* which translates back into the original game as modified payoffs. Ideally this makes the payoffs for defection significantly lower and cooperation spreads far more easily.

In a previous paper [12], we discussed the different ways such as this and other social means that a rational agent might alter their payoffs in a transaction game, such that mutual cooperation becomes an equilibrium for them and most of their opponents. We also discussed the fact that we have to accept that some agents never have a reason to cooperate, such as the malicious agent receiving the irregular payoffs shown in Table 4. An agent like this may be a poster who joins a message board purely in order to push a very specific agenda, so any time their opinion is not being heard loudly they gain no utility from the game. We can see how reputation systems might have very little effect on this agent as a reduction in the payoff for mutual defection is liable to be irrelevant given the strict dominance of defection as a strategy in this game.

Table 4 Malicious Payoffs

Agent/Opponent	Cooperates	Defects
Cooperates	0/0	0/5
Defects	5/0	1/1

An obvious question that arises then is: if we have a global punishment scheme how resistant is a given population to some proportion of interlopers amongst the general population receiving payoffs according to Table 4? We assume that all agents still imitate others who they see as being successful because they individually assume that everyone is playing the same game, and that we have a reputation system which is capable of applying punishment to the extent of reducing a defector's payoff in half at the next time step after defection (with the same noise level as the agents have). The results of this as we vary the percentage of such malicious interlopers are then shown in Figure 6, in which agents interact randomly instead of locally.

Note that those with malicious payoffs are no longer playing a Prisoner's Dilemma game but rather a slightly weighted anti-coordination game, such that the symmetric mixed-strategy yielding optimal payoff is $\frac{3}{8}C + \frac{5}{8}D$ but which with the added punishment phase for defection is reduced back to $\frac{1}{2}C + \frac{1}{2}D$ yielding an expected payoff of $\frac{9}{8}$ which our community still manages to approach.

Likewise in the intermediate stages, if we were to assume that all agents were required to play the same mixed strategy then we could determine that without punishment for 20% interlopers the expected average payoff would be maximal at just over 2.4 with defection occurring in $\frac{5}{96}$ interactions. With punishment the average payoff becomes exactly 2.4 with all agents playing pure cooperation. Over the course of the simulation, we can see that this average payoff is frequently exceeded by virtue of the different types of agents playing individually appropriate strategies. Note that this means the malicious agents are still defecting on a regular basis, but the punishment system is enough to persuade the regular agents to not “stoop to their level” as they would do without it, instead maintaining cooperation the majority of the time.

5 Future Work

Future work could begin by analysing the most successful combinations of strategies which deter defection under the worst conditions of noise and having the global player play this as a mixed strategy against all players. This is likely to allow the relatively lowest level of punishment possible while still maintaining a degree of order. While we require something very closely approaching 100% accuracy in a




Fig. 6 Showing the effects of a varying percentage of interlopers once we have a global punishment system in place. With no global punishment system in place we gain no cooperation.

trading system with tangible resources at stake all we need in a social community is for the overall order to be maintained with no dramatic loss of cooperation.

In application areas where the history of users can be followed by the system, allowing us to guess at the social roles each user is filling at a given point in time, it may make sense to implement a system which encourages the most effective ratio of those behaviours. This could be achieved with a sliding scale of incentives to punish transgressors and/or interact more locally, dependent on the current ratio of cooperation and the existing social roles.

In order to increase the relevance larger population sizes should also be explored, with a focus on adding agents to an existing system, and the rate at which we need to increase local interaction and/or punishments in order to maintain widespread cooperation despite the larger population.

6 Conclusions

We know from Lomborg's earlier results that agents are able to form combinations of strategies which play the iterated Prisoner's Dilemma cooperatively while being protected against invading strategies which might destabilise the system. Here we have examined the circumstances which can lead to the same development of effective global strategies even when agents are playing a far less iterated game.

With the population sizes studied, this is likely to be of more practical use in examining ways in which small communities should grow such that they do not collapse into defection rather than on the global trading systems we would like to focus on in future. In particular, with a recognition that some agents have very little cause to cooperate, it becomes apparent that the efficacy of an externally introduced reputation system does not require it to be strictly effective in punishing defectors in order to achieve the same long term results.

With the emphasis only on preventing a global collapse of trust rather than preventing every instance of defection, approaches which modify the variables of a system in such a way as to encourage the developments seen here are likely to be of use in lower security applications such as online social networks and services which require large scale collaboration without the resources to enforce strong punishments.

References

1. W. T. Au and J. Kwong. *Contemporary Psychological Research on Social Dilemmas*. Cambridge UP, 2004.
2. R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
3. A.L.C. Bazzan, A.P. Cavalheiro. Influence of Social Attachment in a Small-World Network of Agents Playing the Iterated Prisoner's Dilemma. *5th Workshop of Game Theoretic and Decision Theoretic Agents*, 2003
4. K. Binmore, L. Samuelson. Evolutionary Stability in Repeated Games Played by Finite Automata. *Journal of Economic Theory*, 57(2):278-305, 1992.
5. R. Boyd, P.J. Richerson. *Culture and the Evolutionary Process*. University of Chicago Press, Chicago, 1985
6. C. Dellarocas. The digitization of word-of-mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407-1424, 2003.
7. R.I.M. Dunbar. Neocortex size as a constraint on group size in primates. *Journal of Human Evolution* 22:469-493, 1992.
8. Friedman, E. and Resnick, P. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*, 10 (2):173-199, 2001.
9. D. Fudenberg, E. Maskin. The Folk Theorem in Repeated Games with Discounting or with Incomplete Information. *Econometrica*, 54(3):533-554, 1986.
10. D. J.-F. de Quervain, U. Fischbacher, V. Treyer, M. Schellhammer, U. Schnyder, A. Buck, and E. Fehr. The neural basis of altruistic punishment. *Science*, 305:1254-1258, 2004.
11. B. Lomborg. Nucleus and shield: The evolution of social structure in the iterated prisoner's dilemma. *American Sociological Review*, 61(2):278-307, 1996.
12. R. Martin-Hughes, J. Renz. Examining the Motivations of Defection in Large-Scale Open Systems. (to appear) *ACM Symposium on Applied Computing*, 2008.
13. M.E.J. Newman. The Structure and Function of Complex Networks. *SIAM Review*, 45(2):167-256, 2003.
14. A. Rapoport. A Note on the "Index of Cooperation" for Prisoner's Dilemma. *The Journal of Conflict Resolution*, 11(1, Law and Conflict Resolution):100-103, 1967.
15. B. Skyrms. *Evolution of the Social Contract*. Cambridge University Press, 1996.
16. E. Xiao, D. Houser. Emotion Expression in Human Punishment Behavior. *Proceedings of the National Academy of Sciences* Vol. 102(20), pp. 7398-7401, 2005.