Artif Intell Rev (2006) 26:103–114
DOI 10.1007/s10462-007-9039-0

The effects of viscosity in choice and refusal IPD environments

Enda Howley · Colm O'Riordan

Published online: 20 September 2007
© Springer Science+Business Media B.V. 2007

Abstract An objective of multi-agent systems is to build robust intelligent systems capable of existing in complex environments. These systems are often characterised as being uncertain and open to change which make such systems far more difficult to design and understand. Some of this uncertainty and change occurs in open agent environments where agents can freely enter and exit the system. In this paper we will examine this form of population change in a game theoretic setting. These simulations involve studying population change through a number of alternative viscosity models. The simulations will examine two possible trust models. All our simulations will use a simple choice and refusal game environment within which agents may freely choose with which of their peers to interact.

Keywords Trust · Cooperation · Multi-agent systems · Prisoner's dilemma · Implicit trust · Explicit trust

Abbreviations

PD Prisoner's dilemma
IPD Iterated prisoner's Dilemma

1 Introduction

Agent interactions are often heavily biased through certain group structures. These structures are often defined by factors such as trust, reputation, geographical location or kin selection. Within each of these possible scenarios, agent populations rarely remain constant over time. For example, in open multi-agent environments there may be no limitations on agents enter-

E. Howley (✉) · C. O'Riordan
Department of Information Technology, National University Of Ireland, Galway, Ireland
e-mail: enda.howley@nuigalway.ie

C. O'Riordan
e-mail: colm.oriordan@nuigalway.ie

ing or exiting the system. It is therefore important to model this phenomenon in order to understand its effects on agent interactions within the environment. Population viscosity is the measure of change that exists within a population. A highly viscous environment limits movement and displays little or no change over time (Marshall and Rowe 2003). In this paper we will examine the effects of two alternative viscosity models in our population. We will present simulations showing the effects of varying these levels of viscosity on our agent population, in particular, the levels of cooperation and interactions. Our analysis will focus on two trust models and how they differ in coping with the varying degrees of viscosity in the population. We conclude that slight variances in agent trust models can lead to observable differences in agent behaviour when confronted by less viscous environmental conditions. In this paper we hope to address two important questions:

1. What are the effects of reduced population viscosity in an agent environment?
2. Within specific changing environments, what are the effects of the alternative trust models?

In the following section of the paper we present some background research describing previous work in the area of trust and the area of choice and refusal environments. In the Sect. 3 we will outline our simulator design and parameters. The results section will present a series of experiments showing the behaviour of the overall population when alternative forms and levels of viscosity are applied. Stemming from these results we will draw conclusions. Finally, we will discuss the importance of these conclusions within a real world context.

2 Background research

2.1 Choice and refusal models

In game theory there are many commonly used interaction models such as round robin tournaments (Axelrod 1984), spatial (Nowak and May 1993) and tagging techniques (Riolo 1997), (Howley and O'Riordan 2005). These simple models help determine which players interact with each other when playing games such as the Prisoner's Dilemma. Choice and refusal interaction models are an alternative approach to determining player interactions. One choice and refusal model is that presented by Stanley and Ashlock (Stanley et al. 1995). These choice and refusal models allow players to make game offers to possible opponents who are then given the opportunity to accept or reject the game offer. Multi-agent environments very closely mirror this format of choice and refusal, where agents are given the autonomy to mutually agree on their interactions. In multi-agent systems these decisions are often based on notions of trust or reputation.

2.2 Trust

Trust is fundamental to initiating and maintaining cooperative agent interactions in multi-agent environments. Marsh proposes one form of trust classification: He argues that trust may be decomposed into three distinct elements; basic, general and situational trust (Marsh 1994). Basic trust reflects how trusting an agent is towards all its peers. General trust reflects the degree to which an agent trusts each of its peers individually, while situational trust reflects an agent's trust towards a peer, given a number of distinct situations. An agent may be very trustworthy for some relatively unimportant task, while very untrustworthy for a more signif-

Table 1 Payoff matrix

Players choice	Cooperate	Defect
Cooperate	($\lambda 1, \lambda 1$)	($\lambda 2, \lambda 3$)
Defect	($\lambda 3, \lambda 2$)	($\lambda 4, \lambda 4$)

icant one. Furthermore, trust may be communicated or made public which closely resembles reputation models. It can also be represented as a player to player metric in explicit form.

In general, these trust representations involve an agent storing some explicit measure of trust regarding its peers. This can be considered as 'explicit' trust. Regardless of how agents represent levels of trust towards their peers they can express indications of trust implicitly through their interactions.'Implicit' trust involves agents conveying trust through the utilities of their interactions (Howley and O'Riordan 2006). One example from game theory lets players express trust through the payoffs associated with their game interactions. As trust between two players increases the utilities in their Prisoner's Dilemma game change to reflect this.

2.3 The Prisoner's Dilemma

In multi-agent game-theoretic simulations the best known and most commonly used game is the Prisoner's Dilemma (PD). This is a simple two-player game where each player must make a decision to either cooperate (C) or defect (D). Both players decide simultaneously and therefore have no prior knowledge of what the other has decided. If both players cooperate they receive a specific payoff. If both defect they receive a lower payoff. If one cooperates and the other defects then the defector receives the maximum payoff and the cooperator receives the minimum. The payoff matrix outlined in Table 1 demonstrates the potential payoffs for each player.

The dilemma is a non-zero-sum, non-cooperative and simultaneous game. For the dilemma to hold in all cases, two important constraints must be obeyed:

1. $\lambda 2 < \lambda 4 < \lambda 1 < \lambda 3$
2. $2\lambda 1 > \lambda 2 + \lambda 3$

$\lambda 2$ is the sucker's payoff, $\lambda 1$ is the reward for mutual cooperation, $\lambda 4$ is the punishment for mutual defection and $\lambda 3$ provides the incentive or temptation to defect. The dilemma also states $2\lambda 1 > \lambda 2 + \lambda 3$. This constraint prevents players taking alternating turns receiving the sucker's payoff ($\lambda 2$) and the temptation to defect ($\lambda 3$), therefore maximising their score. The following are commonly used values for the Iterated Prisoner's Dilemma $\lambda 1 = 3$, $\lambda 2{=}0$, $\lambda 3 = 5$, $\lambda 4 = 1$. In the non-iterated game, the rational choice is to defect, while in the finitely repeated game, it is rational to defect on the last move and by induction to defect all the time. However, if there exists a non-zero probability the two players will play again, then cooperation may emerge. More extensive background references focusing solely on the Prisoner's Dilemma are presented by Axelrod (1984) and Hoffmann (2000).

2.4 The evolutionary algorithm

In game theory one of the most commonly used means to represent evolutionary simulations are through the use of replicator dynamics. A population consists of n species, each of which adopts a strategy i. The population state can be represented as the following vector at time step t (Generation t):

$$x^t = \left(x^{t_0}, \ldots, x^{t_n}\right) \tag{1}$$

Here x_i^t represents the fraction of the population which can be considered belonging to a species i.

$$\left(x^{t_i} \geq 0, \quad \sum_{i=0}^{n} x^{t_i} = 1\right) \tag{2}$$

The game payoff matrix is used to determine payoff to individual species throughout their lifetime. Payoff to a species i is viewed as the measure of its reproductive success (Smith 1982).

$$|s_i|^t = |s_i|^{t-1} \times \frac{f(s_i)^{t-1}}{\sum_{j=0}^{n} f(s_j)^{t-1}} \tag{3}$$

The representation of a species i in generation t is its representation in generation $t-1$, by the fitness it achieved in generation $t-1$, as a proportion of the average population fitness in generation $t-1$. Hence, the growth rate of an individual species i is proportional to its fitness.

3 Experimental setup

For our experiments we have designed a simulator which allows populations of 125 agents play the Iterated Prisoner's Dilemma over 100 iterations within a simple choice and refusal environment. The simulator uses replicator dynamics as its evolutionary model.

3.1 Choice and refusal

In our simulator the agents offer games in proportion to payoffs received from previous interactions with their peers, i.e., the peers from whom players have received the most payoffs are more likely to be made game offers. When a player is offered a game the decision to accept or reject the game offer is determined probabilistically based on trust. For example, a player is more likely to accept games offered from a trusted peer while less likely to accept games offered from a non-cooperative peer. In the case of an explicit trust model this probability would be determined through levels of average cooperation between the two players involved. In the case of implicit trust this probability would be determined through the game payoffs offered as these are directly related to levels of average cooperation over previous interactions and therefore simply a subtle extension of explicit trust.

3.2 Trust

We model two forms of trust in our simulations. The first is a form of general trust as described by Marsh (1994). We consider this as an explicit trust model as it evaluates trust through tracking each peer's average cooperation $\bar{x}$ to date. Higher levels of average cooperation will result in similar levels of trust. The calculation of average cooperation across n games $\{x_1 \ldots x_n\}$ is represented as follows:

$$\bar{x} = \frac{1}{n} \sum x \tag{4}$$

The second trust model we consider is an implicit trust model which is based on general trust. General trust is normally based solely on average cooperation but this model differs by allowing this trust to influence the temptation to defect in the game payoffs (Howley and O'Riordan 2006). The resulting game remains within the constraints of the Prisoner's Dilemma but the $\lambda 3$ payoff varies linearly to reflect levels of average cooperation (general trust) in each pairwise interaction. This results in the $\lambda 3$ value reflecting levels of average cooperation linearly in the following range.

$$\lambda 1 < \lambda 3 < 2 \times \lambda 1 \tag{5}$$

Since the value of $\lambda 3$ always remains within this specified range all the resulting games comply with the constraints of the Prisoner's Dilemma.

3.3 The strategy set

In order to define a strategy set, we draw upon research by Nowak and Sigmund 1993 (Nowak and Sigmund 1990). A typical strategy includes three primary strategy values representing probabilities of cooperation in an initial move p_i and in response to a cooperation p_c or defection p_d. The resulting strategy genome is as follows:

$$p_i, p_c, p_d, \tag{6}$$

We randomly generate an initial population of strategies using an even distribution to prevent a bias in the initial population towards cooperation or defection.

3.4 Viscosity

In order to simulate change and uncertainty in our agent environment we developed two distinct models of viscosity. Each model probabilistically selects certain members of the agent population for attention. The resulting changes to the population depend on the viscosity model being used. The first model allows viscosity through shuffling. This model introduces a probability of a strategy in the population being selected and shuffled with other selected strategies in the population. Throughout the population these strategies are shuffled while biases between players remain the same. An agent's strategy may yet change through shuffling and agent's peers will still assume its strategy remains the same as previously. The second form of change is viscosity through substitution. This more extreme form of viscosity differs from the previous one in that selected strategies in the population are replaced by new strategies which may not be present in the current population of strategies. These new strategies are selected at random so as not to bias the overall substitution towards more cooperative or non-cooperative strategies. The two viscosity models differ only in the manner they select strategies to replace existing strategies.

4 Experimental results

In this section we present a series of simulations involving the two trust models. We examine the effects of varying certain important environmental conditions involving viscosity in our interaction models. As explained in the previous section, we examine the effects of change throughout the population using two distinct forms of population viscosity. Throughout all our experiments we will examine three important population metrics. First, we will quantify the levels of average cooperation throughout the populations. Second, we will examine the average levels of pairwise interactions throughout the population. In these experiment we will evaluate the levels of connectivity between agents in the population. Agents can be considered nodes in a weighted graph where edges represent interactions between specific nodes in the population. Through calculating the average degree of each vertex we can gain an indication of levels of connectivity throughout the population. If $G = (V, E)$ is a simple graph, and $\deg(v)$ is the degree of a vertex $v \in V$ in G. The average degree in G is defined as follows:

$$D = \sum_{v \in V} \frac{\deg(v)}{|V|} \tag{7}$$

We refer to this metric as average pairwise interactions throughout the population. Finally, we will examine the average levels of repeated interactions which occur throughout the population. In this case each edge has a related weight representing the numbers of interactions between two connected vertices. Let w_{ij} denote the weight of the edge connecting vertex i and j ($w_{ij} = 0$ if there is no edge connecting vertex i and j, and $w_{ij} = n$ if player i and j have had n game interactions). Repeated interactions are calculated using the following equation:

$$W = \sum_{w_{ij} \in V} \frac{w_{ij}}{D} \tag{8}$$

This also gives us a clear indication of how agents interact repeatedly with their most trusted peers. The results presented in this section show averaged results gathered over 20 simulations.

4.1 Reduced population viscosity—through shuffling

In this section we will present a series of simulations examining the effects of reducing population viscosity in the population. In these simulations a random set of strategies were selected and subsequently shuffled in order to simulate some degree of population change.

In the experiment shown in Fig. 4 we examine levels of cooperation when the population is simulated using different levels of population viscosity. We observe the higher levels of cooperation achieved by the explicit trust model over the implicit trust model in totally viscous environments. All levels of shuffling resulted in the implicit trust model converging to a plateau at around 0.5 average cooperation. In the case of the explicit trust model, a small reduction in population viscosity (e.g. 2%) resulted in a fall to about 0.5 average cooperation. But surprisingly further decreases in viscosity resulted in a drop in average cooperation to about 0.4 average cooperation. The implicit trust model showed a greater resilience to higher levels of population change when compared to the explicit trust model in this set of experiments. We believe the resistance of the implicit trust model to change through swapping is due to its ability to promote cooperation in the less cooperative games through a reduced temptation to defect.

In this experiment we examine levels of pairwise interactions. The results shown in Fig. 2 show significant differences in the results achieved by both trust models. We see evidence that the implicit trust model maintains significantly more pairwise interactions throughout the population across multiple generations. We observe the general decline in all pairwise interactions over time as agents bias their interactions away from less trustworthy peers. The explicit trust model shows a greater tendency to limit interactions in early generations while the implicit trust model appears more tolerant in its approach. This difference stems from the

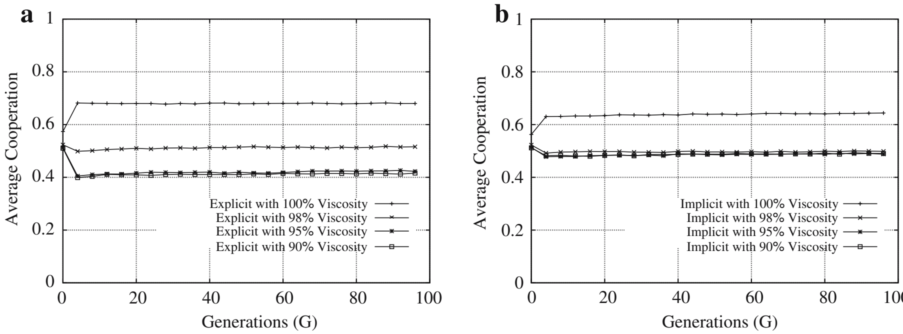

a.Population Using Explicit Trust Model

b.Population Using Implicit Trust Model

Fig. 1 Average cooperation with shuffling population using **(a)** Explicit trust model **(b)** Implicit trust model

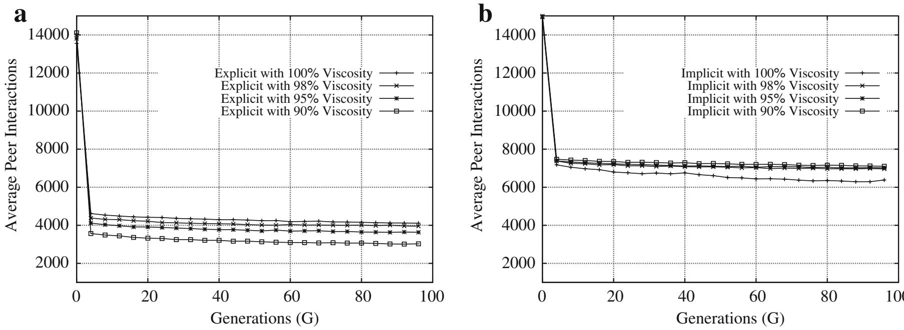

a.Population Using Explicit Trust Model

b.Population Using Implicit Trust Model

Fig. 2 Average pairwise interactions with shuffling population using **(a)** Explicit trust model **(b)**Implicit trust model

implicit trust model tolerating interactions which are intermittently non-cooperative. These interactions tolerate defection when there are high λ3 payoffs available.

In the experiment shown in Fig. 3, we examine levels of repeated agent interactions. We observe the noticeable convergence of the linear implicit trust model to about 60,000 repeated interactions once change through shuffling is introduced into the population. This convergence appears to reinforce our belief that linear implicit can maintain a level of robustness to this form of reduced population viscosity, even at its highest levels.

We believe that these results are directly attributable to the changing payoffs in the implicit trust model. This results in the implicit trust model encouraging cooperation in scenarios where cooperation is infrequent. Increased λ3 values, due to any degree of cooperation, provide an incentive to strategies who cooperate and then avail of heightened λ3 values through defection. Strategies which always defect are heavily penalised in the implicit trust environment as the value of λ3 will always remain low. We notice that the explicit trust model recorded significantly lower levels of repeated pairwise interactions when simulated using increased probabilities of change. Overall the explicit trust model suffers more in a less viscous population.

From the simulations we observe that for environments with lower population viscosity the explicit trust model reacted through decreased numbers of pairwise interactions. This would appear quite rational for increasingly hostile environments but in this case limiting

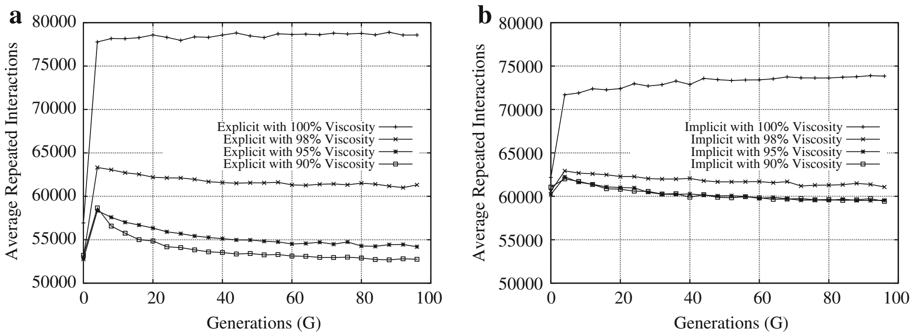

a.Population Using Explicit Trust Model

b.Population Using Implicit Trust Model

Fig. 3 Average repeated interactions with shuffling population using **(a)** Explicit trust model **(b)** Implicit trust model

interactions to a smaller set of trusted peers poses its own specific risks. In highly viscous environments reducing peer interactions has clear benefits through avoiding non-cooperative peers. However, once levels of viscosity fall this behaviour carries a inherent risk. In less viscous changing environments overdependence on a small set of peers poses a greater risk of being exploited by a previously trusted peer. In contrast the implicit trust models tend to interact on average with more peers when in less viscous environments. This stems from its increased chance of interacting with peers who defect intermittently to avail of higher $\lambda 3$ values. The effect of this reduces the risk of over dependance on a small set of peers and thereby reduces the effect of lower population viscosity.

4.2 Reduced population viscosity—through substitution

In this section of our experimental results we examine the effect of introducing the probability of certain strategies in the population being replaced with new ones. The strategies chosen for substitution are selected randomly from an unbiased distribution of strategies.

In the experiment shown in Fig. 4, we measure levels of cooperation when population viscosity is diminished through substitution in our agent population. The explicit and implicit trust models show a decline in average cooperation when viscosity is diminished. The explicit trust model appeared slightly better equipped to cope with reacting to the substitution of non-cooperative strategies. In general, our implicit trust model appeared slower to react to exploitation and rewards cooperation in non-cooperative environments. Alternatively our explicit trust model appeared more direct in classifying which opponents were trustworthy or not. This appears to be beneficial for our implicit trust model in the shuffling experiments but the contrary in this case. In the case of substitution the implicit trust model is heavily exploited through offering high $\lambda 3$ values to invading strategies based on previously very trusting relationships being replaced. Similarly, when cooperative strategies invade previously non-cooperative relationships the relationship links are already cut off from the population and thereby lost. This effect was not nearly as extreme in the shuffling model.

In our final experiments in Figs. 5 and 6 involving substitution we notice the inability of the implicit trust model to limit interactions sufficiently to cope with substitution. While neither model could cope with viscosity through substitution, we can observe that our explicit trust model was slightly better positioned to limit interactions in such a environment.

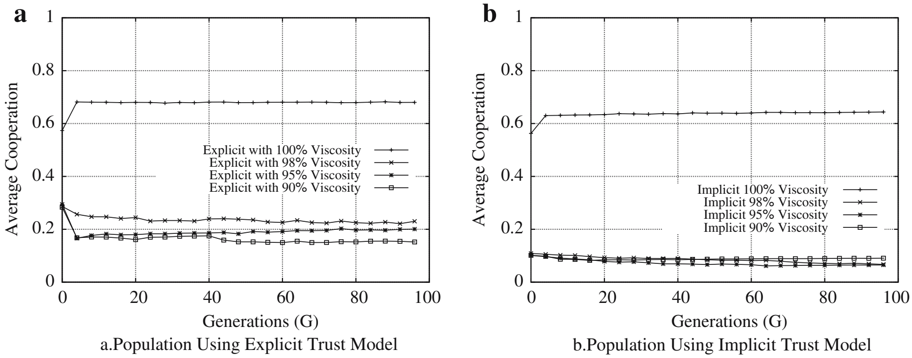

Fig. 4 Average cooperation with substitution population using **(a)** Explicit trust model **(b)** Implicit trust model

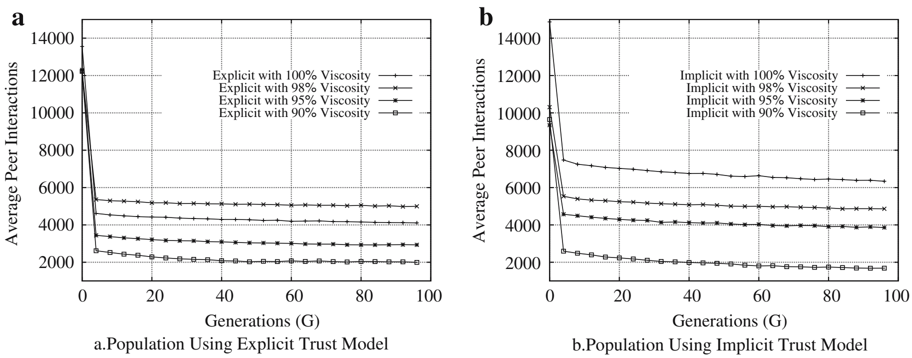

Fig. 5 Average pairwise interactions with substitution population using **(a)** Explicit trust model **(b)** Implicit trust model

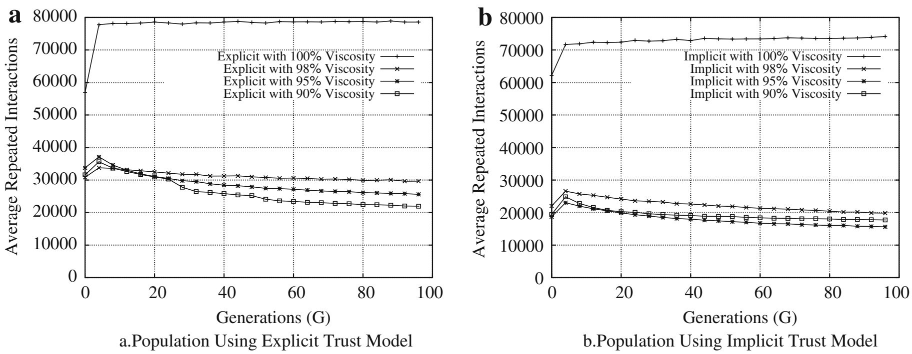

Fig. 6 Average repeated interactions with substitution population using **(a)** Explicit trust model **(b)** Implicit trust model

5 Conclusions

In this paper we have presented a series of simulations involving a choice and refusal game environment where agents bias their interactions using one of two possible trust models. We have studied the effects of introducing uncertainty into the simulation environment using two distinct viscosity methods. This involved modelling reduced viscosity through shuffling or substitution. In all experiments, reduced viscosity resulted in lower levels of cooperation in the population. Change in the population, modelled through shuffling, impacted far less on the population than change modelled through substitution. This difference is due to the likelihood that over time, with shuffling, strategies are more likely to be replaced by similar strategies; this increases the probability of mutually cooperative games. Conversely, substitution increases the probability of introducing strategies that will decrease the level of cooperation in games.

Through analysis of our experiments we identified noticeable differences between the behaviour of both trust models to change within the population. Our implicit trust model proved more capable than the explicit trust model when confronted by shuffling. This stems from the implicit trust model's ability to tolerate minor changes in an opponents behaviour as often occurs in shuffling. Shuffling presents only minor behavioural changes, as the impact of shuffling is dramatically reduced as the population becomes more homogeneous over time. So the initial generations carry the greatest importance when initially biasing interactions. The explicit trust model was less tolerant initially and dramatically cut repeated interactions from any perceived untrustworthy individuals. Alternatively the implicit trust model is by its nature more tolerant and is not as ruthless over initial generations. This offers a clear benefit for later generations when the implicit trust model can benefit from interacting with more individuals repeatedly at a stage when swapping has a reduced impact. In these later generations the explicit trust model has reduced its interactions to a very limited set of peers.

This phenomenon of overdependence presents a number of implications for agent interactions within populations of differing levels of viscosity. Overdependence can prove very risky when in a changing environment of reduced viscosity. This is the first time we have observed through simulation the jeopardy of reducing peer interactions. We conclude that reduced viscosity through swapping has a greater impact on agents with fewer peer interactions as opposed to those with many peer interactions. In the case of reduced viscosity due to substitution we identified the overall failure in both trust models to maintain levels of cooperation. The explicit trust model performed slightly better due to its less tolerant nature but this improvement was marginal.

From a more critical perspective there are some limitations to the experiments presented. In order to simulate open multi-agent environments we propose two possible representations of population viscosity. Many more representations of population viscosity are possible. Another criticism may also include the game and environment definition used. There is no way to exploit the 'beneficial' tolerance of the implicit trust model which we identified. There are undoubtedly strategies which could capitalise on this tolerance.

6 Summary

In this paper we have examined the effects of reduced population viscosity through population shuffling and substitution. Earlier in this paper we posed two important questions. In response to the first question we have observed that lower viscous environments cooperation is more difficult to maintain. It is more difficult to bias agent interactions in less viscous environments as trusted peers are replaced by less cooperative strategies. This is less obvious in our shuffling experiments as over successive generations trusted peers are replaced by strategies who have survived previous generations.

In response to the second question we conclude that there are significant differences between modeling explicit and implicit trust. We have observed that the implicit trust model displays a degree of tolerance which is not present in the explicit trust model. Strategies which are generally cooperative but defect intermittently to avail of high $\lambda 3$ payoffs are tolerated. This behavior is not tolerated in environments which model explicit trust. Finally, we have concluded that implicit trust provides an alternative representation of trust which has not previously been simulated. Through our simulations the evidence indicates that representing trust implicitly provides an inherent form of tolerance with respect to biasing agent interactions.

The results presented in this paper underline the importance of viscosity in agent environments. We have identified how an individuals reliance on a small sets of peers exposes them to the effects of reduced population viscosity. This contradicts the previously understood norms of limiting agent interactions to promote the emergence of cooperation (Howley and O'Riordan 2005).

Previous research involving coalition formation in electronic markets has identified the importance of group size when agents are considering their interaction options (Lerman and Galstyan 2003). In these environments agents begin by joining a small coalition. Larger coalitions are only formed when a number of small coalitions merge.

The study of viscosity is critical to the long term development of robust multi-agent systems. Agent environments are increasingly hostile and challenging through their degree of openness. Electronic markets, for example, are exposed to multiple external stimuli which cannot be predicted or anticipated beforehand. Therefore experiments examining population viscosity are vital to the improved long term design of autonomous agents.

Acknowledgements The primary author would like to acknowledge the Irish Research Council for Science, Engineering and Technology(IRCSET) for their assistance through the Embark initiative.

References

Axelrod R (1984) The evolution of cooperation. Basic Books, New York

Hoffmann R (2000) Twenty years on: The evolution of cooperation revisited. Artif Soci Social Simulat 3(2)

Howley E, O'Riordan C (2005) The emergence of cooperation among agents using simple fixed bias tagging. In: Proceedings of the 2005 congress on evolutionary computation (IEEE CEC'05). vol. 2. IEEE Press, pp 1011–1016

Howley E, O'Riordan C (2006) The effects and evolution of implicit trust in populations playing the iterated prisoner's dilemma. In: Proceedings of the 2006 congress on evolutionary computation (IEEE CEC'06) held as part of the IEEE world congress on computational intelligence (IEEE WCCI'06). IEEE Press

Lerman K, Galstyan A (2003) Agent memory and adaptation in multi-agent systems. In: AAMAS '03: Proceedings of the second international joint conference on autonomous agents and multiagent systems. pp 797–803, ACM Press New York, NY, USA

Marsh S (1994) Trust in distributed artificial intelligence. In: MAAMAW '92: selected papers from the 4th european workshop on modelling autonomous agents in a multi-agent world, artificial social systems. Springer-Verlag, pp 94–112. London, UK

Marshall JA, Rowe JE (2003) Viscous populations and their support for reciprocal cooperation. Artif Life 9(3):327–334

Nowak M, May R (1993) The spatial dilemmas of evolution. Int J Bifurcat Chaos 3:35–78

Nowak M, Sigmund K (1990) The evolution of stochastic strategies in the prisoner's dilemma. Acta Appl Math 20:247–265

Riolo R (1997) The effects and evolution of tag-mediated selection of partners in populations playing the iterated prisoner's dilemma. In: ICGA. pp 378–385

Smith M (1982) The evolution and the theory of games. Cambridge University Press, Cambridge

Stanley EA, Ashlock D, Smucker MD (1995) Iterated prisoner's dilemma with choice and refusal of partners: evolutionary results. In: ECAL. pp 490–502