The Effects and Evolution of Implicit Trust in Populations Playing the Iterated Prisoner's Dilemma

Enda Howley and Colm O'Riordan

Abstract— **The concept of trust is central to engendering cooperation among autonomous agents. This paper focuses on the topic of trust and how agents may bias their interactions based upon implicit trust. We define implicit trust as that which is conveyed through the utilities of a simple game offer. We introduce this concept of implicit trust and present our motivations for examining this phenomenon. We define a game theoretic framework, including possible strategy sets and a game environment. We outline a series of experiments which illustrate the effects of implicit trust. Finally, we draw conclusions based on the experimental results presented.**

I. Introduction

Trust is a fundamental consideration in the development of open, dynamic multi-agent environments. It is an intrinsic key to initiating and maintaining cooperative agent interactions in these challenging environmental conditions. There are relatively few formal definitions of trust, but one chosen by Griffiths and Luck[4] describes trust as a form of risk. They base this partly upon the observations of Marsh that 'entering into a trusting relationship is choosing to take an uncertain path that can lead to either benefit or cost depending on the behaviour of others'[8]. A peer's reliability is determined through their track record, which is represented by a metric called trust. Through classifying risk as an indicator of trust, Griffiths and Luck state that this inverse relationship 'allows us consider the limits of trust more precisely and to quantify it'. They state the relationship between trust T and risk R as the following:

$$R = \frac{1}{T}$$

Griffiths and Luck[4] base their model of trust upon work by Marsh[7] and Gambetta[3]. Stephen Marsh argues that trust may be decomposed into three distinct elements; basic, general and situational trust[7]. Basic trust can be considered as a representation of how trusting an agent is towards all its peers. This can be represented as a single value reflecting a level of trust towards all agents in the environment. This differs from general trust, which reflects how much an agent trusts each of it's peers individually. Here an individual value represents each pairwise matching such as 'X trusts Y'. Situational trust reflects one'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 significant one. Marsh uses the following equation to allow an agent X estimate the situational trust of an agent Y with respect to a situation α. The parameters U, I and T represent the utility, the importance of α and the general trust of X towards Y respectively:

$$T_x(y, \alpha) = U_x(\alpha) \times I_x(\alpha) \times \widehat{T_x(y)}$$

Our proposed implicit representation shares a number of similarities with situational trust. Real world interactions often contain notions of implicit trust; for example, repeatedly buying goods from the same vendor. Once a trusting relationship emerges then the potential stakes can rise reflecting this increased trust. These increased stakes result in a greater temptation to defect which is the critical feature we are hoping to model. From these core principles we hope to examine how agents might successfully bias their interactions among their peers based on implicit trust.

Agent interaction models have been discussed by numerous researchers who have primarily focused on trust and tag based models. However, no clear picture has emerged addressing the known drawbacks of these current techniques. Tagging models have been found to perform badly in highly dynamic environments[6], while also undermining heterogeneity[9].

In order to gain a greater understanding of how agents can successfully bias their interactions we examine an adapted representation of the well known Iterated Prisoner's Dilemma game. This adaptation of the Prisoner's Dilemma is founded on the belief that game interactions which remain static across multiple encounters do not reflect the changing nature of those relationships over time. Previously, changes in these relationships were modelled solely through individual players' strategies. However, an intuitive extension of the game would reflect relationship changes directly through the game parameters. Therefore we define implicit trust as that which is conveyed through the utilities of a simple game offer. In a related article Eriksson and Lindgren proposed and implemented a game theoretic framework involving a repeated game with random observable payoffs[2]. Our proposed framework is outlined in the following sections.

II. The Prisoner's Dilemma

The Prisoner's Dilemma was proposed in 1950 by Melvin Drescher and Merrill Flood, but it was the well known mathematician Albert W. Tucker who coined it's name and

Enda Howley is with the Department of Information Technology, National University of Ireland, Galway, Republic Of Ireland, (email: enda.howley@nuigalway.ie).

Colm O'Riordan is with the Department of Information Technology, National University of Ireland, Galway, Republic Of Ireland, (email: colmor@it.nuigalway.ie).

wrote the first article on the subject[13]. The Prisoner's Dilemma (PD) 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.

TABLE I
PAYOFF MATRIX

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

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:

$$\lambda 2 < \lambda 4 < \lambda 1 < \lambda 3$$

$$2\lambda 1 > \lambda 2 + \lambda 3$$

$\lambda 2$ 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 available by Axelrod[1] and Hoffmann[5].

III. IMPLICIT TRUST GAME

As we have outlined in the previous section the Iterated Prisoner's Dilemma is a simple two player game, within which two players are presented with two options, to cooperate or defect. In this section, we define a new game which follows a similar format. In this game, two players are presented with the same choice, to either cooperate or defect. But the game payoffs presented represent a potential level of risk on the part of the game participants. A rational agent will propose or accept levels of risk based on some indication of trust. One possible indicator of trust would be some function of previous moves such as average cooperation. This concept of risk as an indicator of trust is the underlying principle that motivates this game. As a result the main premise of our proposed game involves modeling implicit trust within a choice or refusal environment similar to that described by Ashlock et al. [12].

There are significant challenges in defining and simulating a strategy space as large as that suggested. We therefore propose a number of specific constraints, which serve to limit the resulting strategy space. The 'temptation to defect' or $\lambda 3$ value changes to reflect previous interactions. This value may be calculated using a various number of methods. We investigated two possible methods of representing this payoff value.

Firstly, we used a simple linear representation of the 'temptation to defect' TD. This is determined using a linear relationship between the explicit metric of 'average cooperation' and the resulting TD. In order to meet the constraints of the PD we enforce a simple rule forbidding this TD value ever equalling $\lambda 1$ or $2 \times \lambda 1$.

Secondly, we examined an equation which is also determined through the explicit metric, 'average cooperation'. This representation allows us to vary the relationship of the 'average cooperation', A and the resulting TD, through using differing ϕ values. Due to the composition of the equation, the Tan function naturally enforces a boundary on the possible range of TD, through the naturally occurring asymptotes.

$$TD = \lambda_l + \frac{\lambda_u - \lambda_l}{2} + \left(\frac{2\tan^{-1}(A \times \phi)}{\pi} \times \frac{\lambda_u - \lambda_l}{2} \right)$$

This calculation for TD represents the new value of $\lambda 3$ (the temptation to defect). The resulting game payoff matrix remains similar to the traditional IPD game.

TABLE II
IMPLICIT TRUST PAYOFF MATRIX

Players Choice	Cooperate	Defect
Cooperate	$(\lambda 1, \lambda 1)$	$(\lambda 2, TD)$
Defect	$(TD, \lambda 2)$	$(\lambda 4, \lambda 4)$

The adapted payoff matrix shown uses the original λ values specified earlier. As with the traditional game these λ values must meet the constraints specified in the traditional IPD game. The value of $\lambda 3$ is calculated using two new values. These define the upper λ_u and lower λ_l bounds of the resulting TD value. This serves to limit the range of TD so it always remains greater than the lower bound λ_l and less than the upper bound λ_u. Since the value of TD may change reflecting the actions of the players it must remain in the following interval range:

$$\lambda 1 < TD < 2 \times \lambda 1$$

Due to this important condition we can simply apply the following rules of thumb for calculating the upper and lower bounds of TD. Throughout our simulations we used this as the standard guide for calculating λ_u and λ_l.

$$\lambda_l = \lambda 1$$

$$\lambda_u = 2 \times \lambda 1$$

Similarly, since TD must remain strictly within these specified bounds we must clearly define the parameter A. This value is a function of previous interactions between the two players and is derived from their average cooperation A to date. The resulting value of A must remain in the range:

$$-1 \leq A \leq +1$$

The temptation to defect value reflects past history through the parameter A in the calculation of TD while remaining within the constraints of the traditional PD game. The values of $\lambda 1, \lambda 2$ and $\lambda 4$ remain constant throughout all player interactions as in the original PD. The value ϕ in our second implicit trust representation is used as a simple scalar quantity through which the relationship of A and TD can be varied. Equally the values of λ_l and λ_u serve to keep all values returned by the function within the acceptable range as represented by $\frac{3\pi}{4}$ and π in the Tan function.

In summary, our game must satisfy the constraints specified in the traditional IPD. These two constraints must be satisfied in order to maintain a Prisoner's Dilemma.

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

As we have previously stated, we have two possible representations of our extended implicit trust game. One uses a simple linear relationship between the parameter A and the resulting TD used. The second uses the following equation to calculate the temptation to defect.

1) $TD = \lambda_l + \frac{\lambda_u - \lambda_l}{2} + \left(\frac{2\tan^{-1}(A \times \phi)}{\pi} \times \frac{\lambda_u - \lambda_l}{2} \right)$

IV. Strategy Set

In order to define a strategy set we draw upon research by Nowak and Sigmund 1993[10]. 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 looks like the following:

$$Genome = p_i, p_c, p_d,$$

Our implicit trust model will make game offers to selected peers through roulette wheel selection. Players using this selection algorithm will construct roulette wheels based on total payoffs received from previous interactions with their peers. As a result of this selection process, peer selection will rapidly progress from initially random game offers to highly structured interactions. The decision to accept or refuse game offers will be determined probabilistically using the game payoffs offered. A higher temptation to defect will represent a higher probability of acceptance. Through successive simulations we will examine the behavior of our implicit trust representation. We will draw comparisons with existing agent interactions models and conclude on the relative merits of implicit trust over these existing models.

V. Environment

Our game environment will simulate a choice or refusal model similar to that proposed by Ashlock et al.[12]. This environment allows agents to propose games, which can then be accepted or rejected. This judgement will be determined probabilistically based on the game offered. Interactions that are initially random will become increasingly more structured as information regarding previous interactions becomes more abundant. The game sequence will follow a series of simple steps:

- Game Offers: Players use roulette wheel selection to make game offers based on previous payoffs received. Choice and refusal decisions are made by players based on the game payoffs offered.
- Game Moves: Agreed games are played over N number of iterations.
- Repeated Game Offers and Moves: Based on repeated game interactions new game payoffs will be calculated and offered.
- New Generation: Each player's fitness is determined through their total game payoffs. Based on these measures of fitness our replicator dynamic will allocate representation in successive generations.

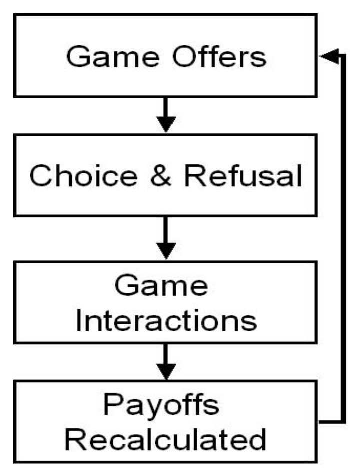

Fig. 1. Game Cycle

We use a simple replicator dynamic in our initial simulations. This evolutionary algorithm simply allocates each player's representation in successive generations based on their fitness.

VI. Experimental Setup

In this section we will present a series of experiments that demonstrate the behavior of our implicit trust models using a series of important metrics. Firstly, we will outline some important experimental parameters. The following is a demonstration of how TD may be calculated using one of our implicit representations. We set the upper and lower bounds of our $\lambda3$ value as λ_u and λ_l. Therefore in our example we use the following λ values.

$$\lambda_l = 3, \lambda2 = 0, \lambda_u = 6, \lambda4 = 1.$$

For simplicity we use a ϕ value of 1. The example shown represents a game pairing where there has been equal levels of cooperation and defection. Therefore average cooperation A is 0.0 when represented on the following scale:

$$-1 \leq A \leq +1$$

Example calculation of TD where $\pi = 180, \lambda_l = 3, \lambda2 = 0, \lambda_u = 6, \lambda4 = 1, A = 0, \phi = 1$:

$$TD = \lambda_l + \frac{\lambda_u - \lambda_l}{2} + \left(\frac{2\tan^{-1}(A \times \phi)}{\pi} \times \frac{\lambda_u - \lambda_l}{2}\right)$$

$$TD = 4.5 + \left(\frac{0}{180} \times 1.5\right)$$

$$TD = 4.5$$

Our experiments are designed to compare the differences and similarities between implicit trust and a simple choice and refusal IPD game. This choice and refusal game will act as a benchmark, through which we hope to assess the behavior of implicit trust. Both implicit and explicit models use identical implementations aside from their ability to reason about which game offers to accept or reject. In this case, our explicit trust simulation uses average cooperation to determine responses to game offers while our alternative model uses implicit trust. The probability of accepting a game offer is determined as follows in our implicit trust and explicit trust models respectively:

$$P_{Accepting} = \frac{TD - \lambda_l}{\lambda_l + \lambda_u}$$

$$P_{Accepting} = A$$

Because implicit trust biases interactions based upon changing payoffs, we must use some fair payoff structure for our IPD model. Here payoffs will remain fixed from game to game but their value must be comparable to those used in the implicit trust model. The fixed payoff matrix used by our simplified choice and refusal IPD is the following:

TABLE III
C/R IPD Payoff Matrix

Players Choice	Cooperate	Defect
Cooperate	$3, 3$	$0, 4.5$
Defect	$4.5, 0$	$1, 1$

Through using the upper and lower bounds of 3 and 6, a temptation to defect value of 4.5 would reflect a game of equal amounts of cooperation or defection. In other words, average cooperation A on a scale of $-1 \leq A \leq +1$ would equal 0.0. In all of the experiments depicted, the length of each game is 100 iterations long. While no two agents can play each other more then 500 times in any given generation. Throughout all our simulations we use a population of 125 agents. This population of agents is initialised using a probabilistically even distribution of strategies. All the experiments shown in this paper use averaged data which was taken over multiple simulations of the same experiments.

A. Average Cooperation

In this experiment we examine the average cooperation of the overall population using alternative agent interactions models. We examine both our implicit trust representations and compare their behavior with a simple choice and refusal model with no ability to bias interactions using implicit trust. As we have stated earlier this simplified model uses a fixed payoff matrix throughout all game interactions while determining choice and refusal using average cooperation. All interaction models use roulette wheel selection based upon previous payoffs received to determine game offers.

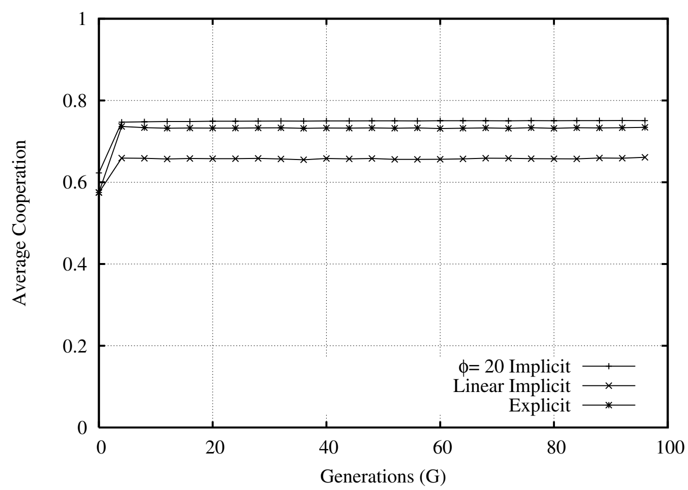

Fig. 2. Average Cooperation

In the experiment shown, we observe the behaviors of our interaction models with respect to average cooperation throughout the population. We observe that each of these interaction models maintain similar levels of cooperation throughout the simulations. The levels of cooperation achieved in each of the simulations is high given that our populations are composed of non-deterministic strategies. We note that both of our implicit models are similar to the explicit trust model in maintaining levels of cooperation throughout successive generations. Across multiple simulations, implicit trust prevented populations of agents from converging to total defection.

B. Average Fitness

In the following experiment we examine a closely related metric involving the overall fitness of the population. We calculate the average fitness of all agents throughout the population based on their average payoffs received in each generation. As in the previous experiment we examine almost identical choice and refusal models. The benchmark interaction model is again used here with a fixed payoff matrix.

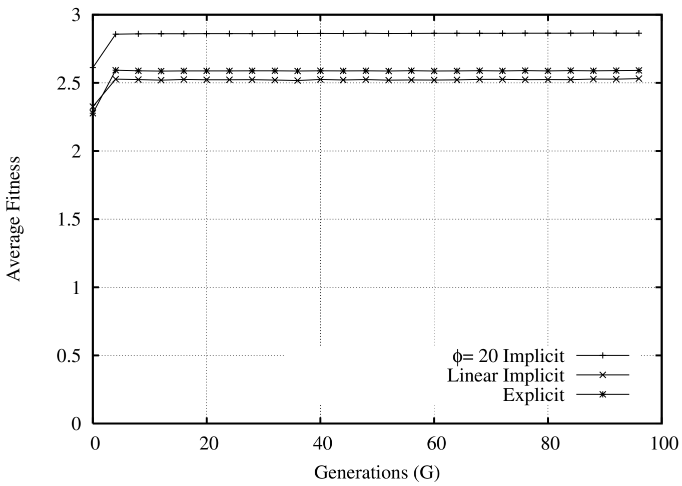

Fig. 3. Average Fitness

In the experiment shown, we observe the behavior differences between our interaction models. We can observe the success of each interaction model in maintaining quite significant levels of fitness throughout the population. We observe the differences in average fitness between our implicit trust models and the explicit trust benchmark. From the experiment shown, we can conclude that the implicit trust models succeed in maintaining high levels of fitness similar to the explicit trust benchmark. We observe that our $\phi = 20$ implicit trust model achieves higher levels of fitness throughout a number of generations. This can be explained through the ability of this implicit trust equation to change TD values more sensitively with respect to previous agent interactions. The resulting changes have a direct and significant effect on how agents bias their interactions.

This similar performance may stem from two possible factors in our implicit models. Firstly, our implicit trust models are successfully biasing agent interactions away from those that are less cooperative. Secondly, due to the changing payoff structure of the implicit trust models, agents may be ‘cashing in’ on heightened TD values which may be built up over a series of cooperative moves. This characteristic of implicit trust is not possible in the explicit model as payoffs always remain static throughout all game interactions. This factor may contribute to the effectiveness of implicit trust in biasing interactions as it directly affects the payoffs received from certain peers who facilitate this ‘cashing in’ phenomenon.

C. Average Number of Peer Interactions

In the following experiment, we examine an important metric which indicates levels of connectivity throughout our population. Through tracking the number of successful pairwise games per agent, we achieve an indication of how our interaction models are successfully biasing interactions towards their peers. As we have seen from previous research, limiting agent interactions contributes significantly to achieving heightened levels of fitness among agents. Tagging mechanisms are extremely successful at boosting agent fitness but these are built upon extreme partitioning of agent populations and high attrition rates in the initial generations[6]. In this experiment, we examine how successfully our model has allowed agents to voluntarily restrict their interactions for their own benefit. Unlike tagging techniques, these interactions are not enforced rigidly by the environment and are instead probabilistically determined through our implicit trust mechanism.

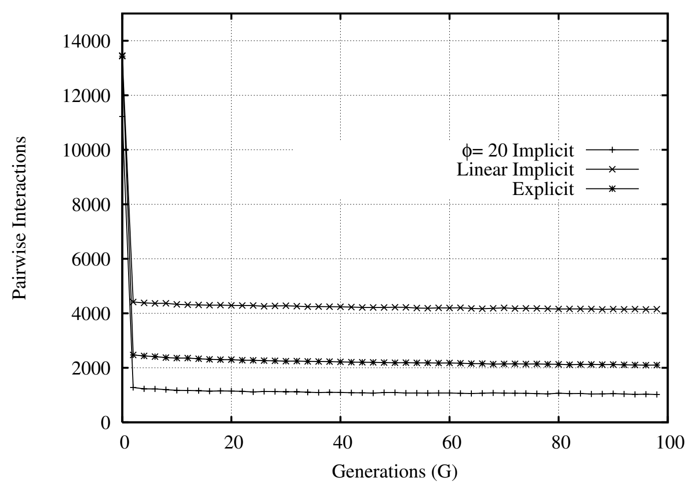

Fig. 4. Average Number Of Peer Interactions

The shown experiment outlines how the respective interaction models limit their interactions over a number of generations. As we can see, our implicit trust models behave similarly to the benchmark through converging to a small limited set of pairwise interactions. This underscores our results in the previous experiments, which showed implicit trust models performing similar to the explicit trust benchmark. We can also identify that models which limited agent interactions more excessively, also achieved the higher degrees of cooperation and fitness. This can be explained partly through a models ability to limit agent interactions to a smaller subset of more reliable peers. The $\phi = 20$ implicit trust model reflected an ability to bias interactions away from untrustworthy peers through its greater sensitivity to change in any players behavior. This contributed heavily to the models behavior and performance in the shown experiment.

D. Average Numbers of Repeated Interactions

In the following experiment, we examine the average number of repeated game interactions per generation. While many agents will choose not to interact for any games, of those who do, we will evaluate the average number of games they decide to play with each other.

From the experiment shown, we can identify the increase in repeated game interactions between peers over successive

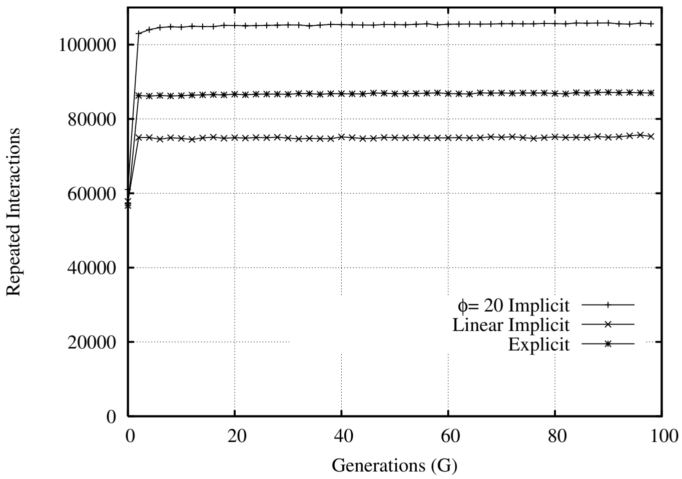

Fig. 5. Average Number Of Repeated Interactions

generations throughout all interaction models. Again our implicit models behave similarly to the explicit trust benchmark model. The implicit trust models show a similar ability to interact repeatedly with selected, more trusted peers quite quickly. The $\phi = 20$ implicit trust model showed a strong ability to bias repeated interactions towards more trusted peers. Such behavior stems from this models heightened sensitivity to change in an opponents behavior. This contributed significantly to the $\phi = 20$ implicit trust model's behavior throughout all our experiments. Participating in many repeated interactions with a limited set of reliable peers offers a distinct advantage to any interaction model and is fundamental to gaining a competitive advantage.

VII. Conclusions

In the previous section, we presented a number of specific experiments. From the results shown, we have identified some of the primary characteristics of implicit trust models. In our experiments, we have demonstrated that agent interactions can be biased successfully through implicit trust. The experiments assessing average levels of cooperation show all interaction models achieving similar high levels of average cooperation and fitness. In the case of both implicit and explicit trust, this is due to their success in biasing agent interactions away from any unreliable peers. The main characteristics of implicit trust are three fold. Firstly, as shown in our experiments our implicit trust models rapidly identify unreliable peers and reduce interactions with them. Secondly, implicit trust significantly increases repeated interactions with peers whom it believes are more reliable and cooperative. Finally we are satisfied that the effect of changing the TD value with respect to previous interactions encourages players to act cooperatively in order to 'cash in' on higher TD payoffs. Because of this phenomenon, players can mutually benefit from successive cooperative gestures, culminating in a defection, to gain the benefit of the increased TD payoff. This is more pronounced in the implicit $\phi = 20$ model as TD can be rapidly increased and availed of by the agents in some alternating manner. The effect of this behavior results in an increase in the average fitness without a noticeable increase in average cooperation throughout the model. This three way benefit contributes significantly to the performance of our implicit trust models.

In the context of previous research, investigating agent interaction models such as tagging schemes, each of our models performed well. While our agent interaction models are significantly different to previous tagging models, our results reflect similar levels of fitness to those identified by Riolo[11]. In his evolved tag bias model, Riolo outlines his hypothesis that, when a population is dominated by mutual cooperation or defection, agents loose their ability to bias interactions based on tags as they are no longer effective for biasing interactions. This remains partially true in the case of our implicit trust model. In some evolutionary environments, where total cooperation or defection is widespread, agents may no longer bias their interactions using specific game offers. We also consider the case where agents are directly affected by the actual game payoffs within each game, in which they choose to participate. This stems primarily from the alternative game payoffs on offer within the implicit trust model. We hope to explore this in more detail throughout future work.

Through proposing this alternative trust model, we hope to have achieved a proof of concept and somehow augmented existing research involving trust in multi-agent systems. Previously proposed models solely use forms of 'explicit trust', which are predominantly strategy level trust metrics, which bias agent interactions regardless of specific utility offers. Situational trust may be the most closely related model to implicit trust but can still be classified as an explicit model, as it is represented exclusively at agent strategy level. Our proposed trust model is a tentative exploration of a new game and strategy space which has not previously been explored. This new agent interaction model draws from a number of areas, such as economic game theory. One of these closely related topics is the area of duopolies. This is the study of corporate trading relationships with respect to splitting of market shares through respective production levels and pricing.

In our future work, we hope to extend our search of our strategy space through more advanced evolutionary approaches. We also hope to conduct more varied simulations of implicit trust using variations of our current definition.

VIII. Summary

In this paper we have proposed the concept of implicit trust. We have outlined our motivations for examining this theory. These stemmed from real world examples and also previous research in the domains of trust and multi-agent systems. We have explicitly defined an interpretation of implicit trust, which included explicitly defining the game constraints. We have outlined a game environment and a possible strategy set from which we have conducted a series of simulations of this game. Through these preliminary experimental results, we have shown that implicit trust can be successfully used to bias agent interactions. We have compared its characteristics with that of explicit trust and identified certain similarities and differences between the models. One of the most notable differences was the emergence of a phenomenon similar to meta-cooperation. Finally we have proposed some future experiments through which we hope to gain a greater understanding of implicit trust and the reasons for its success in the experiments we have described.

Acknowledgment

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

[1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.

[2] A. Eriksson and K. Lindgren. Evolution of strategies in repeated stochastic games with full information of the payoff matrix. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 853–859, San Francisco, California, USA, 7-11 July 2001. Morgan Kaufmann.

[3] D. Gambetta. *Can We Trust Trust?* Basil Blackwell, 1988. Reprinted in electronic edition from Department of Sociology, University of Oxford, chapter 13, pp. 213-237.

[4] N. Griffiths and M. Luck. Cooperative plan selection through trust. In *MAAMAW '99: Proceedings of the 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, pages 162–174, London, UK, 1999. Springer-Verlag.

[5] R. Hoffmann. Twenty years on: The evolution of cooperation revisited. *The Journal of Artificial Societies and Social Simulation*, 3(2), 2000.

[6] E. Howley and C. O'Riordan. The emergence of cooperation among agents using simple fixed bias tagging. In *Proceedings of the 2005 Congress on Evolutionary Computation (IEEE CEC'05)*, volume 2, pages 1011–1016. IEEE Press, 2005.

[7] S. Marsh. Formalising trust as a computational concept. Ph.d thesis, University of Stirling, 1994. http:citeseer.ist.psu.edu/marsh94formalising.html.

[8] S. Marsh. Trust in distributed artificial intelligence. In *MAAMAW '92: Selected papers from the 4th European Workshop on on Modelling Autonomous Agents in a Multi-Agent World, Artificial Social Systems*, pages 94–112, London, UK, 1994. Springer-Verlag.

[9] A. McDonald and S. Sen. The success and failure of tag-mediated evolution of cooperation. In *LAMAS*, 2005.

[10] M. Nowak and K. Sigmund. The evolution of stochastic strategies in the prisoner's dilemma. *Acta Applicandae Mathematicae*, 20:247–265, 1990.

[11] R. Riolo. The effects and evolution of tag-mediated selection of partners in populations playing the iterated prisoner's dilemma. In *ICGA*, pages 378–385, 1997.

[12] E. A. Stanley, D. Ashlock, and M. D. Smucker. Iterated prisoner's dilemma with choice and refusal of partners: Evolutionary results. In *ECAL*, pages 490–502, 1995.

[13] A. W. Tucker. A two-person dilemma. Stanford University, Discussion Paper, 1950.