The Effects of Payoff Preferences on Agent Tolerance

Enda Howley and Colm O'Riordan

Department of Information Technology
National University of Ireland, Galway
enda.howley@nuigalway.ie
colm.oriordan@nuigalway.ie

Abstract

An objective of multi-agent systems is to build robust intelligent systems capable of existing in complex environments. These environments are often open, noisy and subject to rapid, unpredictable changes. This paper will explore how agents can bias their interactions and choices in these complex environments. Existing research has investigated how agents can bias their interactions based on factors such as similarity, trust or reputation. Unfortunately, much of this research has ignored how agents are influenced by their preferences for certain game payoffs. This paper will show that individual payoff preferences have a significant effect on the behaviors that emerge within an agent environment. We argue that agents must not only determine with whom to interact, but also the levels of benefit or risk these interactions should represent. This paper presents a series of game theoretic simulations examining the effects of agent payoff preferences within an evolutionary setting. Our experiments show that these factors promote tolerance throughout the population. We provide an experimental benchmark using an almost identical game environment where payoffs are not considered by agents. Furthermore, we also present simulations involving noise, thereby demonstrating the ability of these more tolerant agents to cope with uncertainty in their environment.

Introduction

Agent interactions are often heavily biased through certain group structures. These structures are often defined by factors such as geographical location (Axelrod, 1984), kin selection (Hamilton, 1963), choice and refusal (Stanley et al., 1995), or trust and reputation (Dellarocas, 2003). In the real world, individuals often bias their interactions through factors such as their need for certain services, or preferences for particular goals. As a result, we must acknowledge that not all agent interactions are identical and are driven by individual preferences and needs. Therefore in this paper we will examine through game theoretic simulations how agent payoff preferences influence the overall agent population. Some researchers have examined the payoffs commonly used in the Prisoners Dilemma and concluded that certain payoffs promote cooperation (Fogel, 1993). The implications of agent payoff preferences when determining their peer interactions remain to be fully explored and understood.

This paper shows how game payoff preferences directly influence the levels of tolerance and reciprocity throughout an agent population. Existing research has not examined the significance of these agent payoff preferences. For example, in a multi-agent environment an agent may only trust one of its peers. Yet, in order to satisfy its individual preferences this agent may choose to interact with a less trusted peer. This decision could be based on a payoff preference or similarly a preference for a service being offered. In short, within a game theoretic model, agent interactions should reflect that agents are free to bias their interactions based on their preferred peers and also their preferred games. As a result we will simulate an environment where agents may offer and choose games based upon their preferred payoff values. Some agents are risk takers and prefer games which have higher risk payoffs, while others are more risk averse and prefer game payoffs which hold lower risk.

Previous research of this type has focused primarily on IPD games which remain constant throughout the population. Recent research has started to explore the effects of allowing the game payoffs to change (Taylor and Nowak, 2006; Howley and O'Riordan, 2006a, 2008). The need to study these interactions stems from the fact that real world interactions rarely remain identical indefinitely. In reality, an agents interactions will be determined through its bias towards preferred individuals and its bias towards achieving specific goals. This paper will investigate aspects of this statement and in particular address the following two research questions:

1. What are the effects of game payoff preferences on the overall agent population and the strategy genes?

2. What strategies are most successful in this variable payoff environment when noise is introduced?

In the following section of the paper we outline our motivations and aims. We will then discuss various aspects of background and related research from the subject domain. This will involve existing research in the areas of spatial, tagging and trust models. Subsequent sections will describe our evolutionary algorithm, simulator design, parameters and experimental setup. The results section will present a series of experiments showing the behavior of the overall population when alternative interaction models are used. Stemming from these results we will outline our conclusions.

Aims and Motivations

The primary goal of this paper is to propose a simple extension to the traditional Iterated Prisoner's Dilemma and analyse the effects of this salient extension on agent interactions. Previous research from the authors has examined a number of agent interaction models. In previous work, we examined the underlying dynamics of tag mediated interaction models. We identified the significance of tag group size on the levels of cooperation that emerge within the agent population. We also examined the effects of openness on trust between agents (Howley and O'Riordan, 2006b). We found that alternative forms of openness and change effected agents in very different ways. We also identified the dependencies that can emerge between agents within a trust model when exposed to openness. More recent research has investigated payoff variances and preferences among agents (Howley and O'Riordan, 2008). We outline how agents like to avail of higher payoffs through defection but this results in the agent population exploiting each other into choosing games with low risk payoffs. These strategy preferences emerged to dominate the agent populations.

This paper examines the agent strategies at gene level and presents a more detailed examination of their preferences for certain game payoffs. Furthermore we hope to clarify the heightened levels of tolerance present in these game environments. We hope to ascertain the scale of these differences and the reasons behind them. Also we hope to test these levels of tolerance through introducing noise into the game interactions. This differs from simulating reduced population viscosity, or mutation which would serve to undermine cooperation in cooperative groups. In this work we are more interested in identifying the strategy traits that emerge when these factors are not present.

We have extended the Prisoner's Dilemma to reflect variable payoffs, thereby scientifically capturing the effects of agent payoff/risk preferences. Agents' who express their game preference based on a games 'temptation to defect' are in effect specifying a unique game with an associated degree of risk. We provide a more detailed description of this extension in later sections.

Through the results presented in this paper, we aim to extend our previous research while also demonstrating the important dimensional space which has largely been ignored by existing research in multi-agent systems. The differences shown in this paper present many implications for the domain of multi-agent research. The most important of these involves the need to delineate between agent environments where all interactions are of equal value and those where interactions are not equal.

Background Research

In this paper our main topic of concern involves how agents bias their interactions. Previous research on this topic has examined techniques such as spatial, tagging, kin selection and trust. In this section we will discuss some of the existing research on these topics. We will also introduce and discuss the Iterated Prisoner's Dilemma (IPD), which is used throughout our simulations.

Spatial, Tagging and Kin Selection

In relation to the emergence of cooperation, one of the most important considerations involves how agents bias their interactions towards cooperative peers and away from non-cooperative peers. In this paper we are only concerned with the latter. Kin selection is one such interaction mechanism involving groups of related individuals (Hamilton, 1963). Another more common interaction model involves agents located on a spatial topology such as a grid (Axelrod, 1984; Nowak and May, 1993). Agents bias their interactions and therefore play peers located on adjacent cells of the grid. Tag-mediated interaction models are based on a similar premise. These models locate agents on an abstract topology and then bias interactions based on players proximity to each other (Holland, 1993; Riolo, 1997). Arbitrary tags are similar to visible markings or labels which may be used by agents to bias their interactions based on their preferences. Some real world examples of tags could include football fans recognising each other from their jerseys or travelers recognising each other abroad through their native accents. Tags can provide a more general representation of agent interactions than spatial models. Later in this paper we will outline how tag-mediated selection may be used to structure interactions based on players individual preferences for certain games.

Iterated Prisoner's Dilemma

The Prisoner's Dilemma (PD) is a simple two player, non-zero sum, non-cooperative game. 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.

For this game to be classified as a dilemma in all cases, certain constraints must be adhered to. The following is the first constraint:

$$\lambda 2 < \lambda 4 < \lambda 1 < \lambda 3 \qquad (1)$$

Table 1: Prisoner's Dilemma Payoff Matrix

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

These conditions result in $\lambda 2$ being 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 second constraint is the following:

$$2\lambda 1 > \lambda 2 + \lambda 3 \tag{2}$$

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 λ values are commonly used in the 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. This game has been used throughout numerous research domains, including economics, computer science and the social sciences. More detailed discussions on the Prisoner's Dilemma and its various guises are widely available (Axelrod, 1984; Hoffmann, 2000; Delahaye et al., 2000; Kendall et al., 2006).

Evolutionary Algorithm

The experimental results presented in this paper involve agent environments simulated over successive generations. In each of these generations an evolutionary algorithm is applied to reflect the real world pressures on under performing entities and alternatively reward the best performing ones. In this section we will outline in detail the evolutionary algorithm used throughout our simulations.

In the domain of game theory, one of the most common evolutionary techniques involves replicator dynamics. These quite general evolutionary models replicate changes in agent's fitness through increasing or decreasing their representation in successive generations. Therefore in a population 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):

$$x^{\ t} = (x^{\ t_0}, ..., x^{\ t_n}) \tag{3}$$

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

$$(x^{\ t_i} \geq 0, \sum_{i=0}^{n} x^{\ t_i} = 1) \tag{4}$$

The game payoffs represented in the payoff matrix are used to determine payoff to individual species throughout their lifetime. Payoff to a species i is viewed as an indicator of fitness and thereby a measure of its reproductive success (Smith, 1982).

$$js_ij^t = js_ij^{t-1} \times \frac{f(s_i)^{t-1}}{\sum_{j=0}^{n} f(s_j)^{t-1}} \tag{5}$$

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.

Uncertainty and Noise

This paper presents a number of noise experiments which investigate the effects of uncertainty on the agent population. In previous research when we first examined payoff preferences among agents, we observed significant levels of intermittent defections (Howley and O'Riordan, 2008). Because of this we undertook this more detailed investigation of these population dynamics which appeared to promote forgiveness among agent strategies. In addition this paper also examines this phenomenon through simulating noise in the game environment.

Existing research has simulated uncertainty through a series of methods involving noise (Bendor et al., 1991) and reduced population viscosity (Howley and O'Riordan, 2006b). In this paper we have used noise as a means of simulating uncertainty. This will serve to test the ability of agents to forgive opponents. Previous research has shown that more forgiving and generous strategies perform best in noisy environments (Bendor et al., 1991).

Simulator Design

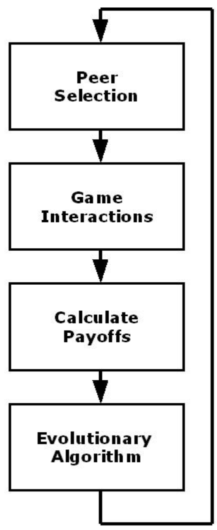

Figure 1: Game Cycle

In this section we outline our overall simulator design. We begin with an introduction to the game cycle (Figure. 1). We describe how we extended the Iterated Prisoners Dilemma (IPD) to allow agents express preferences for certain types of games. We also outline our strategy set.

Firstly, agents play their selected opponents. They enter IPD games using payoffs acceptable to their individual payoff preferences. The payoff preference of the offering agent us used and the opponent can then chose to interact or not. This decision is based on its own payoff preference. Once all games are played game payoffs are totaled and averaged. These are then used as measures of fitness for our evolutionary algorithm. This evolutionary algorithm uses replicator dynamics based on fitness to determine agent representation over successive generations. No crossover or mutation is applied.

Agents select their opponents probabilistically based on the proximity of their tags. This is done in a similar way to much previous research involving tag mediated interactions. Players with similar tags are far more likely to interact than other pairings where there may be some difference between players tags. This differs from common green beard dynamics, whereby by individuals identify each other through their beard colour and cooperate if their beards are a similar colour and defect if they are different (Dawkins, 1976). Our model limits peer interactions to individuals of a similar tag value, and makes no assumption about an individuals actions towards those of a specific tag value.

Strategy Set

In order to define a strategy set, we refer to existing research which uses three bit IPD strategies (Nowak and Sigmund, 1990). In our simulations each strategy genome includes four genes representing probabilities of cooperation in an initial move p_i, in response to a cooperation p_c and defection p_d. The final strategy gene p_t represents an individuals game payoff preference. Some strategies are be more inclined to prefer lower risk games while others will prefer higher risk games. This is similar to people who are often natural risk takers while others are more risk averse. As we will explain in the following section our game environment permits players to agree the ‘temptation to defect’ (TD) value in the Prisoner’s Dilemma game. The resulting strategy genome looks like the following:

$$Genome = p_i, p_c, p_d, p_t \tag{6}$$

The Variable Payoff Prisoner’s Dilemma

The extended IPD game remains similar to the original game described earlier. It remains a simple two player dilemma which is non-zero-sum, non-cooperative and played simultaneously. For this game to remain a Prisoner’s Dilemma it must still remain within the constraints of the original game as mentioned earlier. This game differs in that the payoffs used in each game interaction are not always the same. The extended game uses the following adapted payoff matrix. In this game the $\lambda1$, $\lambda2$, $\lambda4$ payoffs remain constant while in this extended game the value of TD is determined by the individual players involved in each game interaction.

Table 2: Adapted IPD Payoff Matrix

Players Choice	Cooperate	Defect
Cooperate	$(\lambda1, \lambda1)$	$(\lambda2, TD)$
Defect	$(TD, \lambda2)$	$(\lambda4, \lambda4)$

For this game to remain a valid IPD, then the value of TD must remain within the following range of values:

$$\lambda1 < TD < 2 \times \lambda1 \tag{7}$$

The IPD payoff values used throughout this research are as follows: $\lambda1 = 5000$, $\lambda2 = 0$, $\lambda3 = TD$, $\lambda4 = 1$. As stated above the value of TD must always remain within the following range: $\lambda1 < TD < 2 \times \lambda1$. These λ values provide an expressive range of possible TD values.

Our decision not to allow agents determine all game payoffs stems firstly from the need to maintain a valid Prisoner’s Dilemma, and secondly that all interaction choices be based on a fair and equal footing. One can also argue that a Prisoner’s Dilemma which allows the TD to change is identical to a bounded Prisoner’s Dilemma where all payoffs are permitted to change but still remain bounded by an upper payoff limit. This is due to all the payoffs being interdependent and relative as specified by the PD constraints. Therefore by allowing the TD payoff to change, all the game payoffs change relative to each other.

Experimental Results

In this section, we present a series of simulations involving our multi-agent population. We present direct comparisons between a number of multi-agent environments when using fixed payoff games versus variable payoff games. Firstly we examine these differences under noiseless environmental conditions. Subsequently, we present this comparison using a noisy game environment, whereby agent actions are effected by a degree of noise which will demonstrate more clearly the emergence of tolerance in our variable payoff game simulations.

All the simulations outlined in this paper involve populations of 1000 agents. Each experiment is a aggregation of 50 experimental runs. Each game interaction lasts 20 iterations.

The first interaction model is a variable payoff model where agents agree a TD payoff depending on their respective p_t genes. As in a tag environment, players choose their peers based on their tag (p_t gene) similarity. In this model the p_t gene value reflects a players preferences for games of a certain value. A high p_t gene would result in a high TD payoff game while a low p_t gene would result in a low TD payoff game. This results in players with similar p_t genes interacting in games with TD payoffs proportional to their p_t genes. To determine the probability of two individuals interacting we use a formula proposed in previous research on tag-mediated interactions (Riolo, 1997). The dissimilarity of two individuals (A and C) is defined as follows:

$$d_{A,C} = |Ap_t - Cp_t| \tag{8}$$

The second interaction model is almost exactly the same. This is a fixed payoff model which allows the players determine their peer interactions based on their p_t gene similarity. The one significant difference is that all games use the same fixed payoffs. $\lambda 1 = 5000$, $\lambda 2 = 0$, $\lambda 3 = 7500$ and $\lambda 4 = 1$.

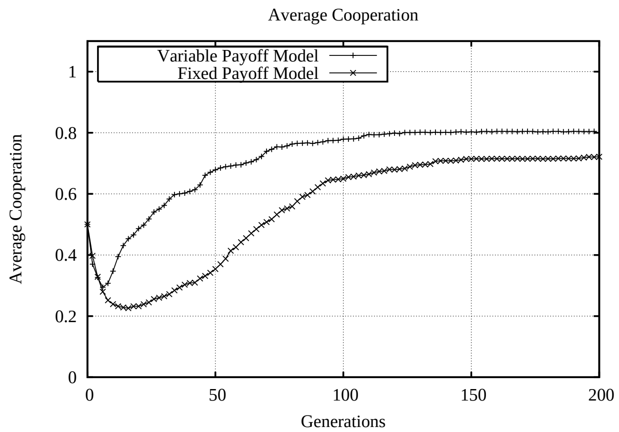

Figure 2: Average Cooperation.

We observe in the data shown in Fig. 2. the levels of cooperation attained using two agent models. These levels of cooperation were significantly higher in the variable payoff model. Furthermore these levels of heightened cooperation were also much faster to emerge. The table below presents a number of statistics reinforcing these observations. On average this model remained 0.10 greater than the static model. The data shown in Table 3 indicates the scale of the differences between the fixed and variable payoff environments. These differences were found to be statistically significant.

Model	μ	σ
Fixed	0.6102	0.1664
Variable	0.7103	0.1390

Table 3: Average Cooperation in Noiseless Environment

The following experiments show the average values for each strategy gene respectively. These values represent averages taken throughout the agent population at the start of each generation. From the results shown, we can ascertain the levels of reciprocity and tolerance present throughout the agent population. These can be identified from examining the p_c and p_d gene values respectively.

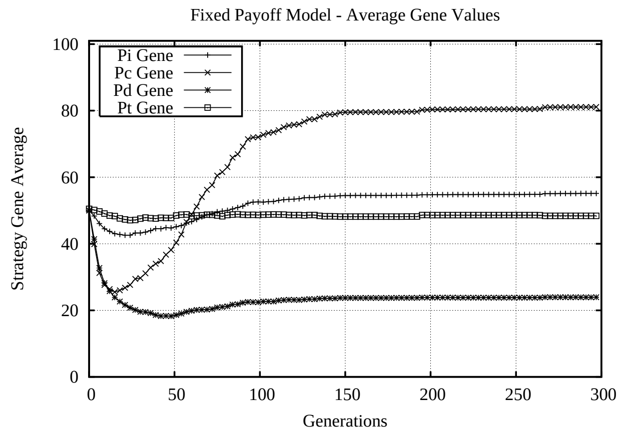

Figure 3: Fixed Payoff Game - Strategy Genes.

Fig. 3 shows average gene values recorded in the static payoff game environment. We observe how the p_t gene remains almost static. This gene experiences no evolutionary pressures as it serves simply as a tag for biasing interactions. The value of this gene is completely random from experiment to experiment. Therefore it's mean value across a large number of experiments always remains close to 0.5. The levels of cooperation in this model as shown in the first experiment (Fig. 2) can be attributed to the significant increase in the p_c gene from generation 20 onwards.

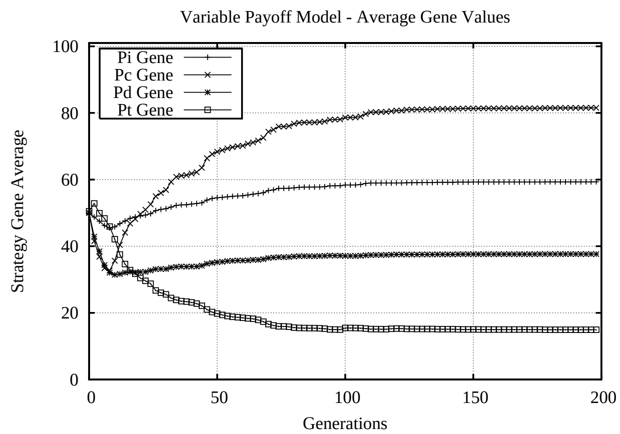

Figure 4: Variable Payoff Game - Strategy Genes.

The levels of cooperation identified in the variable payoff model in the initial experiment (Fig. 2.) are justified through the data shown in the in Fig. 4. This graph shows the p_d gene reaching levels that are significantly higher than in the static payoff model (Fig. 3.). As a result, agents are more likely to cooperate after a defection. This indicates a degree of tolerance or forgiveness which is far greater than in the static payoff environment. Furthermore, the average p_c gene values indicate a similar likelihood of cooperation following an opponents cooperation. This characteristic emerges very rapidly in the initial generations of the variable payoff environment and is associated with a increased likelihood of mutual cooperation in an agent environment.

Summary The results shown here demonstrate the clear differences between static and variable payoff game environments. The heightened levels of cooperation identified in the variable payoff environment have been shown to emerge through specific genetic differences in that strategies that perform best in the respective game environments. These cooperation levels stem from the emergence of tolerance among the participant strategies. Defection in the static payoff model exerts no evolutionary pressure on the agent game preferences. Yet defection in the variable payoff model exerts pressure on agents to chose game interactions with lower TD game payoffs. This results in an agent population who are predominantly cooperative and also possess low p_t genes. Therefore any subsequent defections in this population incur few penalties and are tolerated throughout the agent population. These conclusions are confirmed through our analysis of the strategy genes in each game environment. The average p_t gene value initially increases as strategies choose high TD payoff games. They avail of these high TD payoffs by exploiting their peers, and this rapidly becomes more common throughout the population. Subsequently, these strategies start to mutually defect and begin to suffer. The p_t gene levels fall dramatically as cooperative groups emerge to dominate the population. Strategies who intermittently exploit to avail of certain TD payoffs thrive in this environment. These are the underlying reasons behind the increased tolerance and generosity throughout the variable payoff model.

Noisy Environments In the previous section we examined the similarities and differences between our static and variable payoff game environments. In order to more rigorously test our explanation of the differences between the two game environments, we will now examine their respective dynamics under noisy conditions. We represent noise as a probability that a move will be inverted from C to D or vice versa. The following experiments show the levels of cooperation recorded when alternative degrees of noise are simulated.

Fig. 5. shows the levels of cooperation involving simulations using fixed payoff games and alternative levels of environmental noise. From the simulations shown we can clearly see the effects of noise on levels of cooperation. As would be expected, 1% noise has a noticeable effect on cooperation while 5% has a much more dramatic effect throughout. These results show the extent to which strategies in the fixed payoff environment can cope with intermittent defections. High levels of tolerance would be very beneficial to individuals hoping to cope with intermittent defections.

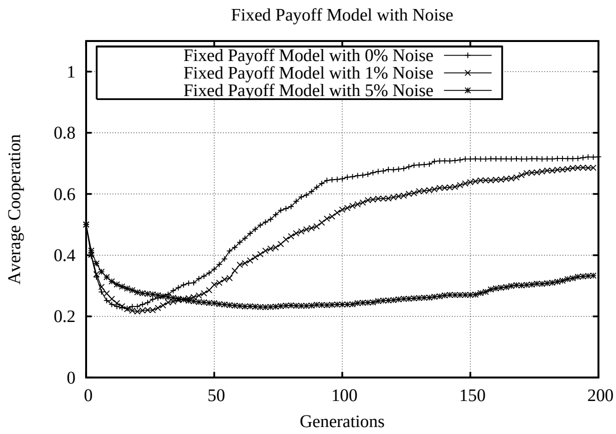

Figure 5: Average Cooperation

Model	Noise	μ	σ
Fixed	0%	0.6102	0.1664
Fixed	1%	0.4810	0.1661
Fixed	5%	0.2718	0.0390
Variable	0%	0.7103	0.1390
Variable	1%	0.6579	0.1396
Variable	5%	0.4924	0.1042

Table 4: Average Cooperation in Noisy Environments

Fig. 6. shows the effects of noise on levels of cooperation in a variable payoff environment. The main differences with the previous experiment in Fig. 5. are the levels of cooperation recorded for 5% noise. The strategies in the variable payoff environment appear to cope much better to these levels of noise. This is reinforces our earlier conclusions that these variable payoff environments promote tolerance throughout an agent population.

The final set of graphs show how the gene strategy values evolved within two game environments when 5% noise was introduced. The simulations shown in Fig. 7. represent the fixed payoff environment. These results show the same non convergence of the p_t gene, as its value carries no great significance in the fixed payoff model. The p_i and p_d genes fall in value while the p_c gene is the only gene which appears to recover in spite of the noise. The slow convergence of the p_c gene continues for about another 300 generations and reaches a level slightly below that identified in the noiseless experiment shown in Fig. 3. More significantly are the values of of the p_d gene which remain very low and indicate low levels of tolerance throughout the agent population. This contributes strongly to the levels of cooperation identified in Fig. 5. for this game environment using 5% noise. It is clear that any occurrences of intermittent defections as

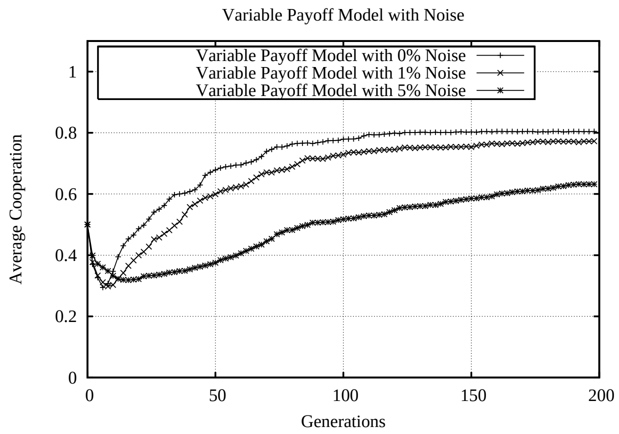

Figure 6: Average Cooperation

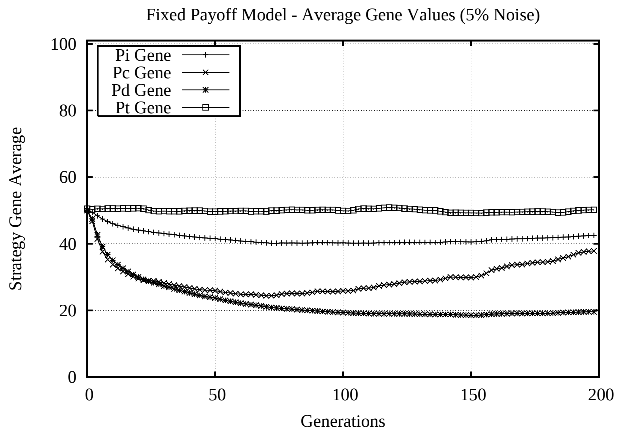

Figure 7: Average Cooperation

a result of noise will result in mutual defection between the participating individual agents.

Fig. 8. shows the gene values recorded in the variable payoff model with 5% noise. This experiment shows higher levels of tolerance as shown through the p_d gene values. Each of the strategy genes in this model continue to converge in a similar way to their noiseless counterpart presented in Fig. 4. but at a much slower rate. These indicate a much higher level of tolerance throughout the game environment and directly contribute to the heightened levels of cooperation identified previously. They explain the fundamental underlying dynamics which resulted to the significant results presented in Fig. 2.

Summary This section has examined a series of specific experimental results. These have shown the underlying differences between static payoff game interactions and the alternative agreed variable payoff model. These differences stem from the increased ability of strategies in the variable payoff model to forgive and tolerate intermittent defections.

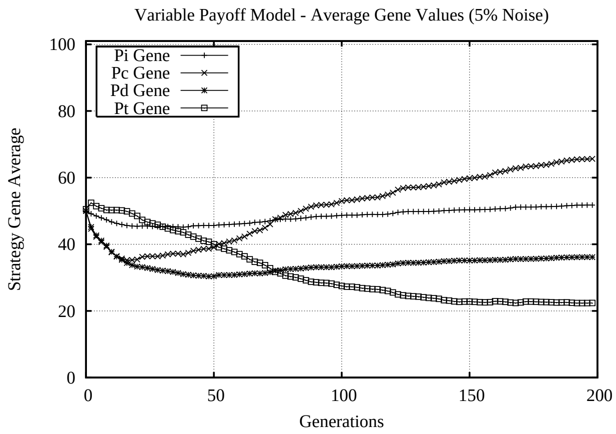

Figure 8: Average Cooperation

This is shown to be even more significant in a noisy environment where intermittent defections are more common. The statistical analysis shown in the tables indicate the significance of these differences. By demonstrating the each environment's ability to cope with intermittent defections, these noise experiments show the main factors that explain the differences first identified in Fig. 2. These noise experiments are not intended as an extensive examination of noise and its effects on IPD strategies. That would involve simulating many more levels and forms of noise.

Conclusions

This paper has presented two alternative game theoretic environments where agents play the Iterated Prisoner's Dilemma. Players bias their interactions through using a designated strategy gene. Through a series of experiments, we have shown that agent behavours can be fundamentally effected by the introduction of a variable payoff game. To date most research has been based on static payoff games and the resulting conclusions have been adopted and cited by many multi-agent researchers. We argue that variable payoff games provide a more realistic basis for real world agent interactions.

Initially we posed two fundamental research questions. Firstly, we queried the influence of variable payoffs on the agent environment. We have shown that this extension resulted in a significant increase in the numbers of strategies using higher value p_d genes. These resulted in greater levels of forgiveness throughout the agent population. We have also shown as previously (Howley and O'Riordan, 2008) that these agents favor lower value p_t genes. This indicates their preference for lower TD payoff games and reinforces the reasons why these strategies are more tolerant of defections. Reduced payoff rewards for defections would naturally make the agent more tolerant of such non cooperation. Initial exploitation for high TD games provides a significant advantage to strategies who only play games involving lower risk games. These strategies then thrive and dominate the population.

Our second question queries which are the most successful strategies. From analysing the trends throughout all the experiments, it is clear that the strategies that are successful in fixed payoff environments are not successful in variable payoff environments and vice-versa. The most successful strategies in the fixed payoff environments are highly reciprocal and therefore not very tolerant. The most successful strategies in the variable payoff environment are highly tolerant and prefer low TD games. Once noise is introduced in the variable payoff environment, the strategies that are tolerant and encourage cooperation perform the best.

This paper has shown the fundamental differences between fixed and variable payoff environments from both a high level analysis and also a gene level examination. We have shown the intrinsic ability of variable game environments to encourage cooperation through tolerance. This leads to higher degrees of cooperation in both noiseless and noisy environments. These differences show the importance to future researchers, of differentiating between multi-agent environments where all agent interactions hold identical significance, and those which offer alternative rewards. This presents researchers with the possibility of encouraging tolerance throughout an agent population without making any assumptions about the agent population.

Acknowledgements

The primary author would like to acknowledge the support of the Irish Research Council for Science, Engineering and Technology for their support through the Embark Initiative.

References

Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York.

Bendor, J., Kramer, R., and Stout, S. (1991). When in doubt cooperation in a noisy prisoner's dilemma. *Journal of Conflict Resolution*, 35:691–719.

Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press, New York.

Delahaye, J. P., Mathieu, P., and Beaufils, B. (2000). *The Iterated Lift Dilemma: how to establish meta-cooperation with your opponent*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.

Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Manage. Sci.*, 49(10):1407–1424.

Fogel, D. B. (1993). Evolving behaviors in the iterated prisoner's dilemma. *Evol. Comput.*, 1(1):77–97.

Hamilton, W. (1963). The evolution of altruistic behaviour. *The American Naturalist*, 97:354–6.

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

Holland, J. (1993). The effects of labels (tags) on social interactions. *Working Paper, Santa Fe Institute 93-10-064*.

Howley, E. and O'Riordan, C. (2006a). 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.

Howley, E. and O'Riordan, C. (2006b). The effects of viscosity in choice and refusal ipd environments. *Artif. Intell. Rev.*, 26(1-2):103–114.

Howley, E. and O'Riordan, C. (2008). Agent interactions and implicit trust in ipd environments. In Tuyls, K., Nowe, A., Guessoum, Z., and Kudenko, D., editors, *Adaptive Agents and Multi-Agent Systems III*, volume 4865 of *Lecture Notes in Artificial Intelligence*, pages 87–101. Springer Verlag, Berlin.

Kendall, G., Yao, X., and Chong, S. Y., editors (2006). *The Iterated Prisoner's Dilemma: 20 Years On*. World Scientific.

Nowak, M. and May, R. (1993). The spatial dilemmas of evolution. *Int Journal of Bifurcation and Chaos*, 3:35–78.

Nowak, M. and Sigmund, K. (1990). The evolution of stochastic strategies in the prisoner's dilemma. *Acta Applicandae Mathematicae*, 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*, pages 378–385.

Smith, M. (1982). *The Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

Stanley, E. A., Ashlock, D., and Smucker, M. D. (1995). Iterated prisoner's dilemma with choice and refusal of partners: Evolutionary results. In *Proceedings of the Third European Conference on Advances in Artificial Life*, pages 490–502, London, UK. Springer-Verlag.

Taylor, C. and Nowak, M. (2006). Evolutionary game dynamics with non-uniform interaction rates. *Theor Popul Biol*, 69:243–252.