Games With Dynamic Payoffs

Shashi Mittal, Sumeet Agarwal
mshashi@iitk.ac.in, sagarwal@iitk.ac.in

Department of Computer Science and Engineering,
Indian Institute of Technology,
Kanpur, India.

April 23, 2005

Abstract

In this paper, we present a study on games in which the payoffs are dynamic, that is the payoffs change with time, or are dependent on the strategy profiles of the players. We investigate two such games: a Language acquisition game, and the Spatial prisoner's dilemma game. While games in which the payoffs are constant have received much attention, games with dynamic payoffs have not been studied extensively in literature. We show in this paper that how such games can model real life situations closely, like acquisition of a grammar by children and evolution of culture.

1 Introduction

Game theory is a fascinating topic in economics and mathematics. Of late, game theory has found a lot of practical applications in the area of computer sciences, particularly resource allocation and selfish network routing. In this paper, we use the concepts from game theory to model two natural phenomena : One is of language acquisition, and the other of spatial prisoner's dilemma.

The paper is organized as follows : First, we briefly discuss the basic concepts of games and game theory. Then we define games with dynamic payoffs, and how these games are different from the usual games. In the next two sections, we describe the language acquisition game, and the spatial prisoner's dilemma game. In these two sections, we present the model employed, the simulations carried out and the results obtained from the simulations. Finally we discuss the significance for these results, and how these models can be extended in future work.

2 Basics of game theory

Game theory is essentially the study of multi-agent decision problems. A **game** is a situation, where more than one agents are competing for limited payoffs. Each agent can make one among several moves available to him, and the payoff to each agent depends upon the moves taken by the other agents in the game.

A game has the following components :

1. **A set of Players** : This is denoted by $P = \{P_1, \ldots, P_n\}$, where n is the number of players in the game.
2. **Set of rules** : The rules specify how the game is to proceed.
3. **Set of strategies** : The set of strategies of player P_i are denoted by $S_i = \{S_{i1}, \ldots, S_{is_i}\}$ where s_i is the number of strategies player P_i has.
4. **Set of outcomes** : O, the possible strategy combinations of all the players in the game.
5. **Payoffs** : $u_i(o)$ for each player P_i and each outcome $o \in O$.

For example, consider the game of Prisoner's Dilemma shown in figure 1. For this game, there are two players, Player 1 and Player 2. The set of strategies of each player is the same, $S = \{C, D\}$ (C for cooperate, D for defect). There are certain rules which each player has to follow, and he can safely assume that others are following these rules. In this game, each player can choose either cooperate of defect. The player has to act independently of the other player, and both have to choose their move simultaneously. The payoffs for each player corresponding to each strategy combination is

		Player 2	
	Strategies	Cooperate(C)	Defect (D)
Player 1	Cooperate (C)	(3, 3)	(0, 5)
	Defect (D)	(5, 0)	(1, 1)

Figure 1: The classical choice for payoff in Prisoner's Dilemma (row player payoffs are given first)

given in the figure. We will discuss the spatial version of this game in detail in later part of this paper.

The strategy profile of a player can either be *pure* (in which case he always makes the same move), or it can be *mixed*. A mixed strategy of a player P_i is a probability distribution $\{p_{i1}, \ldots, p_{is_i}\}$ over the set of his strategies, such that

$$\sum_{j=1}^{s_i} p_{ij} = 1 \quad (1)$$

for all players $P_i \in P$. For example, in the game of prisoner's dilemma, a mixed strategy for the two players can be $\{0.4,\ 0.6\}$ and $\{0.5,\ 0.5\}$. This means that, if the game was played repeatedly for 100 moves, then P_1 would choose C about 40 times, and D around 60 times, whereas P_2 would choose C and D about 50 times each. The payoff for each player can be calculated using the probability distribution. In this case, for example, the payoff of P_1 will be $0.4 \times 0.5 \times 3\ +\ 0.4 \times 0.5 \times 0\ +\ 0.6 \times 0.5 \times 5\ +\ 0.6 \times 0.5 \times 1$, which is 2.4. Similarly, the payoff of the second player comes out to be 1.9.

Games can either be *zero sum* (or *constant sum*, which means that the sum of payoffs of all the players for all strategy combinations is zero), or *non zero sum*, that is, the sum of payoffs of all the players is not necessarily the same for all strategy combinations. For example, the iterated prisoner's dilemma game is a zero sum game. A game of tic tac toe is a zero sum game.

More details on games and their applications can be found in (Neumann & Morgenstern, 1947), (Osborne & Rubenstein, 1994).

3 Games with dynamic payoffs

In the iterated prisoner's dilemma, the payoffs of the players for a strategy combination is a constant. In games with dynamic payoffs, the payoffs are not fixed, and can change with time (as in the case of spatial prisoner's dilemma), or can be dependent on the strategy profile of the players (as in the language acquisition game). Such games are more difficult to analyze, but are useful in modeling many naturally occurring phenomenon.

4 The language acquisition game

One of the most important aspects of languages in humans in the grammar. A grammar is a set of rules that maps linguistic form and meaning. There are many interesting features of language acquisition in humans. One is called the *poverty of stimulus*, that is, the evidence available to the child does not uniquely determine the underlying grammatical rules. Nevertheless, children reliably achieve correct grammatical competence, at a very young age. This is known as *the paradox of language acquisition*.

According to Chomsky (Chomsky, 1965), a child needs a pre-formed linguistic theory that can specify candidate grammars that might be compatible with the available linguistic data. This pre-specified theory is called the *Universal Grammar*. Hence, for language acquisition, the child needs a mechanism for processing input sentences and a search space of candidate grammars from which to choose an appropriate grammar. Here, since we are studying language acquisition mainly in the context of acquiring the correct grammar, therefore we will use the terms "language acquisition" and "grammar acquisition" synonymously. Also, we will be using the terms "child", "player" and "agent" interchangeably.

A detailed exposition on various models of language acquisition can be found in (Niyogi, 2004). Jain et al (Jain et al., 1999) present various learning models, that can be applied to language acquisition as well. Here, we present a game theoretic model to explain how language acquisition takes place in humans. The language acquisition game is inspired by the model proposed by Komarova et al (Komarova et al., 2001), (Komorava & Nowak, 2002) to explain the population dynamics of grammar acquisition. It has also been employed to explain the natural selection of the critical period of language acquisition (Komarova & Nowak, 2001).

4.1 Description of the game

This game can be describes as follows :

- **Agents** : The game can have Kagents, represented as $A_1, \ldots, A_K$. These agents are trying to acquire a grammar by interacting with other agents. In this paper, we focus our discussions only for the case where $K = 2$.

- **Set of strategies** : Each agent has the same strategy set, given by $G = \{G_1, \ldots, G_n\}$. Here, G represents a *Universal Grammar Set*, i.e. the set of grammars from which the agent has to choose a grammar for himself.

- **Strategy profile** : The strategy profile of agent A_i is given by $s_i = \{p_{i1}, p_{i2}, \ldots, p_{in}\}$ where p_{ij} = probability that A_i uses G_i for communicating with other agents. Note that $\sum_{j=1}^{n} p_{ij} = 1$.

- **Rules** : The 2 agent language acquisition game is simulated as follows :

  1. A move in the game corresponds to the two agents speaking out a sentence simultaneously. The agent A_i speaks the sentence using grammar G_{ij} with probability p_{ij}.
  2. Each agent gets to know the sentence spoken by the other agent in the previous move. The agents checks this sentence against each grammar of the universal grammar set, and on the basis of it updates his probability of using each of the grammars.

  The similarity between grammars is modeled using a similarity matrix, $\{s_{ij}\}$. s_{ij} represents the probability that a sentence spoken using grammar G_i will be compatible with the grammarG_j. Komarova et al have proposed a method by which this similarity matrix can be computed.

- **Payoffs** : Let x_i = fraction of agents using the grammar G_i. Then,

$$x_i = \frac{1}{K} \sum_{j=1}^{n} p_{ji} \tag{2}$$

where K = number of agents in the population (in this paper, we consider only $K = 2$). For each grammar G_i, we define the fitness corresponding to the grammar f_i as follows :

$$f_i = \frac{1}{2} \sum_{j=1}^{n} x_j (s_{ij} + s_{ji}) \tag{3}$$

The payoff of the agent A_i is given by the equation

$$\psi_i = \sum_{j=1}^{n} f_j p_{ij} \tag{4}$$

We further make a simplification by assuming a *fully symmetrical system*, in which $s_{ii} = 1$ for all $i = 1, \ldots, n$ and $s_{ij} = s$ for $i \neq j$, $0 < s \leq 1$. For such a system, in a 2 agent game, the respective payoffs of the players, ψ_1 and ψ_2 are

$$\psi_1 = s + \frac{1}{2}(1 - s) \sum_{j=1}^{n} p_{1j}(p_{1j} + p_{2j}) \tag{5}$$

$$\psi_2 = s + \frac{1}{2}(1 - s) \sum_{j=1}^{n} p_{2j}(p_{1j} + p_{2j}) \tag{6}$$

Note that the maximum possible payoff of any agent cannot exceed 1, and it cannot be less than s. We assume that the agents are rational, that is, each agent is trying to maximize his payoff. However, none of the agents has any idea about the strategy (i.e. the probability set) of the other agent(s).

4.2 Learning mechanism

We employ the following learning mechanism for each agent : After hearing a sentence σ, the agent A_i checks if σ is compatible with G_j. If the sentence is compatible, then the probability p_{ij} is updated as follows:

$$p_{ij}^{new} = p_{ij}^{old} + \delta \tag{7}$$

where δ is a parameter which can either be fixed, or can change with the number of sentences heard. After comparing σ with all the grammars in the universal grammar set, the probabilities for the agent a re-normalized (so that $\sum_{j=1}^{n} p_{ij} = 1$).

If a fixed value of δ is used, then it is observed that the probability values for the two agents change in a random fashion, and no equilibrium is obtained as such. Therefore we employ a simulated annealing learning model, in which δ is given by

$$\delta = 1 - e^{\frac{-(1-\psi)}{kt}} \tag{8}$$

where t = number of sentences heard by the agent, ψ is the payoff for that particular agent, and k is a variable that can be changed (the effect of k parameter on grammar acquisition dynamics is described later in the paper). Such a function for δ satisfies the following two important properties

1. The value of δ decreases with increase in t, that is, the change in probabilities of using grammars decreases with time.

2. The value of δ is smaller for higher values of ψ, that is, higher the payoff for that agent, lower will be his tendency to change to a different probability profile.

4.3 Simulations and results

We investigated the outcome of the above game for 2 players with small number of grammars ($n = 3$) in the universal grammar set. The results of the simulations are described here.

In general, the results can be summed up as follows : The probabilities to which the two agents converge to depends upon the initial conditions. For low values of s (i.e., when the similarity between the grammars in low), the two agents tend to attain the same probabilities of using the grammars. For higher values of s, the probabilities remain close to the initial values, and no agreement between the two agent emerges. If a higher value of k is used, the probabilities tend to reach a steady value quickly, although the two agents do not tend to converge to the same probabilities for the grammars. If a low value of k is used, the probabilities tend to fluctuate arbitrarily. For all the simulations described here, we have used a value of 1.00 for k, which gives an optimal performance with respect to learning for both the agents.

Figures 2 and 3 show how the probabilities change for the two agents for $s = 0.0$ (i.e. no similarity between the grammars of the universal grammar set) and a given set of initial probabilities. In this case, a mixture of two grammars emerges as the equilibrium between the two agents. This shows that the agents are indeed able to converge together on the usage of particular grammar(s). Figures 4 and 5 show how the fitness of the two agents vary with time. It is clear that the two agents are able to reach a good fitness level after speaking only a few sentences.

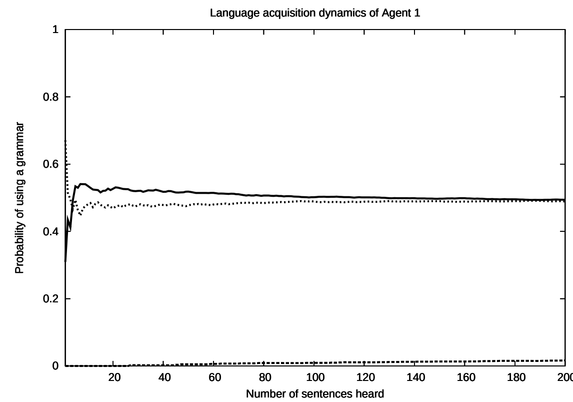

Figure 2: Plot of strategy profile of Agent 1 for $s = 0.0$

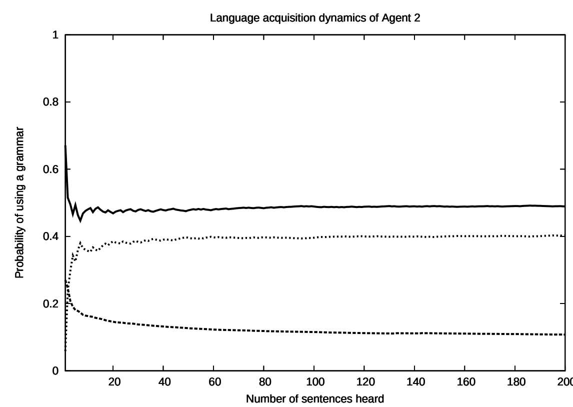

Figure 3: Plot of strategy profile of Agent 2 for $s = 0.0$

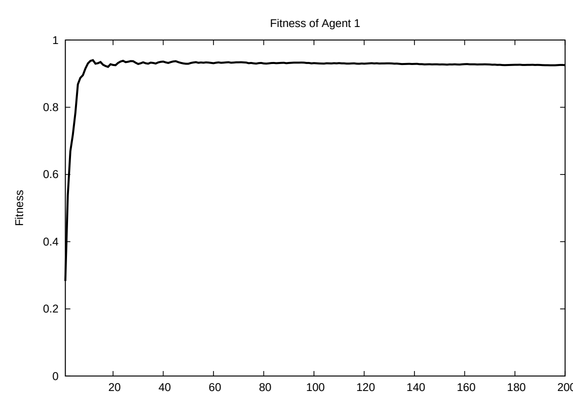

Figure 4: Plot of fitness (payoff) of Agent 1

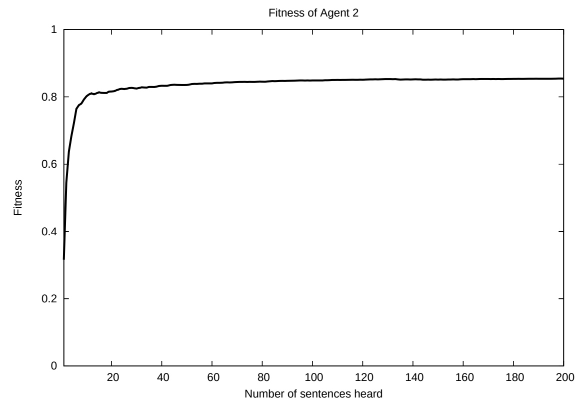

Figure 5: Plot of fitness (payoff) of Agent 2

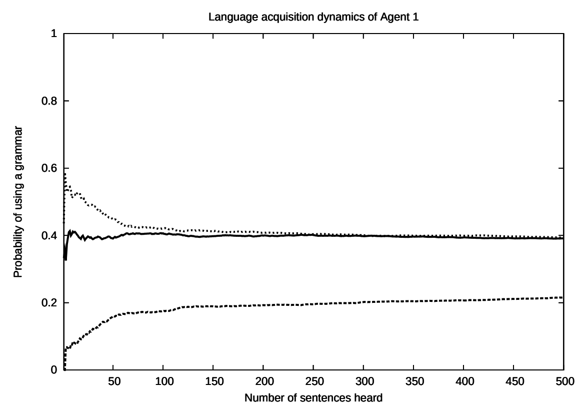

Figure 6: Plot of strategy profile of Agent 1 for $s = 0.25$

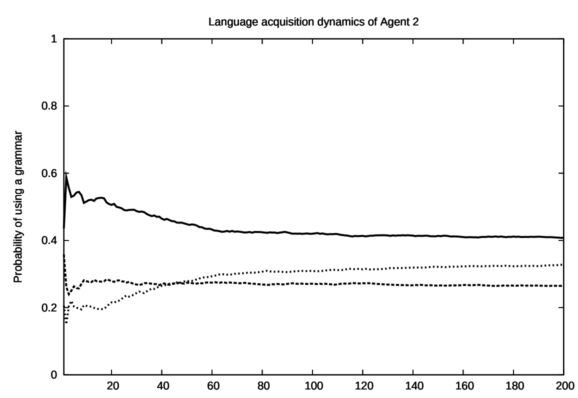

Figure 7: Plot of strategy profile of Agent 2 for $s = 0.25$

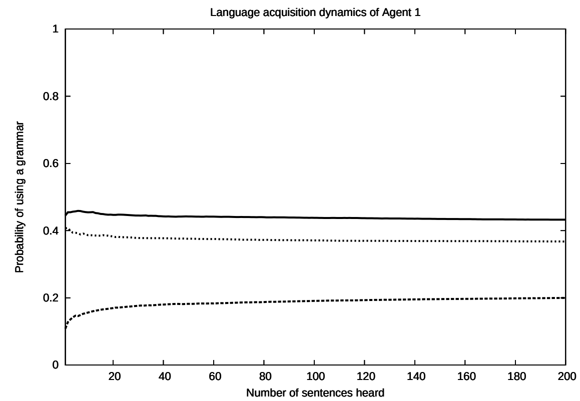

Figure 8: Plot of strategy profile of Agent 1 for $s = 0.80$

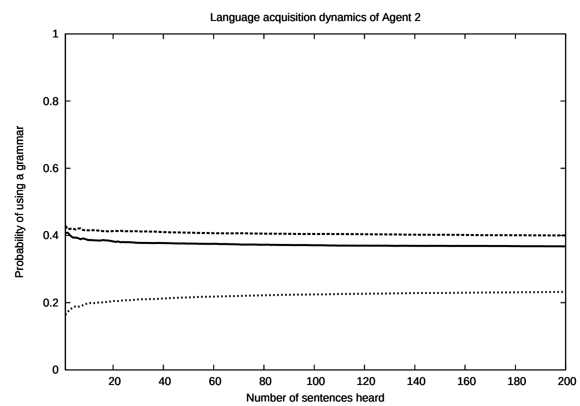

Figure 9: Plot of strategy profile of Agent 2 for $s = 0.80$

When $s = 0.25$, the two agents still tend to attain the same probabilities of using the grammars, however this time the convergence of the two agents to the same probabilities is not as sharp as it was in the previous case. This case is shown in figures 6 and 7. With increasing value of s, the two agents do not tend to converge to same grammars. This is because the similarity between the grammars increases, and therefore they can communicate successfully with each other even if they use different grammars. As an example, the plots for the case $s = 0.80$ have been shown in figures 8 and 9. The plots clearly reflect the trend mentioned before. When $s = 1.00$, the probabilities do not change at all, i.e. probabilities values remain same as the initial values throughout. This is expected, since in this case, the two agents can always communicate successfully irrespective of the grammars they use.

4.4 Evaluation of the model

The model for language acquisition proposed here works remarkably well for two person case. Though such cases, where two individuals are simultaneously trying to acquire a grammar without any external help have not been empirically studied, nevertheless this model makes an attempt to show how the two agents will acquire a grammar simultaneously. It would be interesting to study the dynamics of grammar acquisition when there are more than two agents in the population, and with more grammars in the universal grammar set.

5 The Spatial Prisoner's Dilemma

The most famous of all the situations modeled using game theory is the classic Prisoner's Dilemma. In the basic version of this game, we have two prisoners being interrogated by the police. Both have the option of either confessing their crime or denying it. The police offer them a bargain whereby if one of them confesses and the other doesn't, the one who confessed goes scot-free, while the other gets 5 years in jail. If both confess, they get 3 years each, but if neither does so, then the police has insufficient evidence to convict them, so they only get 1 year each on a lighter charge. This situation is represented by the payoff matrix in figure 1. For a single game, the dominant strategy (which is sub-optimal) is to confess (or defect, represented as D in the figure). However, the game becomes more interesting when played in an iterated fashion, with the opportunity to respond to the opponent's moves.

The spatial prisoner's dilemma game extends the concept to a two-dimensional grid. Each individual is surrounded by 8 neighbors, and plays the game with each of them simultaneously. On any given turn, a player must adopt one common strategy for all the games, and then will receive separate payoffs from each one (depending on the particular opponent's strategy), as per the payoff matrix. The total payoff will be defined as the sum of these 8 individual game payoffs. After each turn, a player may switch strategies based on his result in the previous round. Tomochi and Kono (Tomochi & Kono, 2002) studied this game using a copycat learning algorithm, whereby each player would look at the payoffs of all his neighbors and himself, and then pick up the strategy of the one with the highest payoff among these for the next round.

	C	D
C	R, R	S, T
D	T, S	P, P

Figure 10: Payoff matrix for Spatial Prisoner's Dilemma

The game is made more interesting (and realistic) if we allow for changes in the payoff matrix. For instance, we may think of this game as modeling social interactions in general, where cooperative behavior is more desirable for the general good, but competitive behavior may bring higher individual gain. To offset the tendency to behave competitively, society might therefore impose a punishment on such behavior, and this punishment might be increased with increase in the number of competitors, in order to try and push individuals back towards cooperation. With this general idea as the basis, Tomochi and Komo (Tomochi & Kono, 2002) proposed an evolution rule for the payoff matrix. The general form of this matrix can be represented as in Figure 10.

Here, S < P < R < T. Now, we may fix S=0 and T=1 as the endpoints of our payoff scale, and vary only P and R, always subject to the condition P < R. The following equations were used for this purpose:

$$P(t+1) = P(t) - kg(p_D(t) - p_D^*) \tag{9}$$

$$R(t+1) = R(t) + k^{'}g(p_D(t) - p_D^*) \tag{10}$$

Here, g is some function, k and $k^{'}$ are constants, $p_D(t)$ is the proportion of defectors in the population at time t, and p_D^* is the tolerance limit, i.e. if defection level surpasses it, then punishment for defection is correspondingly increased. For our simulations, we set g(x) = x; k = $k^{'}$ = 0.01; and $p_D^* = 0.5$. In addition to the aforementioned copycat method, we also ran simulations using the two most successful strategies in the two-player prisoner's dilemma, tit-for-tat (TFT) and Pavlov, in

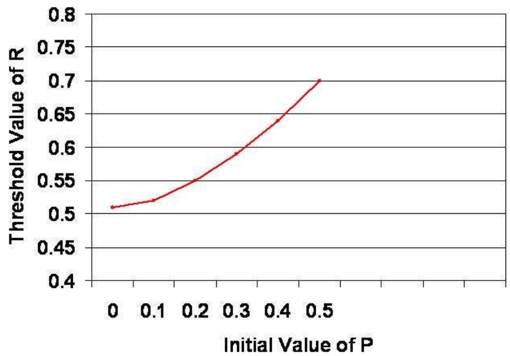

Figure 11: Threshold Values of R for Emergence of Cooperation with Copycat Updation

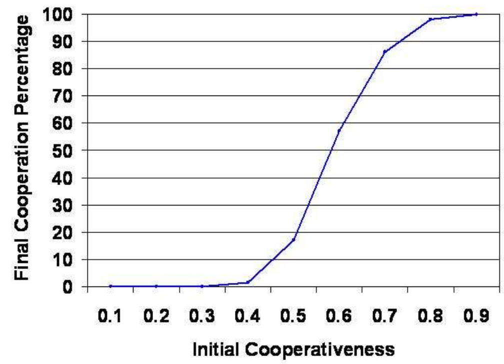

Figure 12: Stable Cooperation Levels with TFT Updation

order to see how these different methods of updation influenced the evolution of cooperation in the population. The results of these simulations are reported in the next section.

5.1 Method

The simulations were run in Java, using separate matrices to represent the member strategies and payoffs. TFT was simulated by having each member, at each turn, adopt the strategy that was adopted by the majority of his 8 neighbors on the previous turn. In case of a 4-4 split, the individual would retain his old strategy. Pavlov was simulated by having each individual switch strategies if and only if his payoff for the previous turn was less than 4 (half of 8, the maximum possible payoff). Parameters like the initial values of P and R and the initial cooperation level were kept fixed in some cases and varied in others. After the initial cooperation level had been fixed, the distribution of initial strategies ('C' or 'D') among the members was randomly done. Techniques and results for each of the three updation methods are now presented in more detail.

5.2 Copycat Updation

For copycat updation, the technique used was that on each turn, each individual looked at the payoffs of all his neighbors, as well as his own payoff (on the previous turn), and from amongst these 9 options chose the strategy (on the previous turn) of the one with the highest payoff, as his own strategy for the coming turn. A few preliminary simulations showed that the final convergent cooperation level and final average payoff level were essentially independent of the initial cooperation level. Thus, the initial cooperation level was fixed at 0.5 for all the simulations, i.e. half of the population started with strategy 'C', the other half with strategy 'D'. Grid size was fixed at 1000x1000, and 200 rounds were played per simulation. It was seen that for a given initial value of P, there existed a threshold initial value of R, i.e. if the initial value of R was above the threshold, cooperation would evolve, otherwise it would not. So we focused our efforts on finding these threshold initial values of R for different initial values of P (from 0.1 to 0.5). The results are shown in Figure 11.

In all cases when R was initially below the threshold, the final cooperation level was 0, as was the final average payoff. When R started out above the threshold, cooperation level converged to around 50 percent (varying from 35-65), while average payoff converged to around 4.3 (varying from 3.5-5.0). As can be seen from the figure, the threshold initial value of R rises with rise in the initial value of P, going from around 0.5 ($P(0) = 0$) to 0.7 ($P(0) = 0.5$). Thus, we see that the final state of the population is entirely dependent on the initial values of P and R, and is independent of the initial strategy distribution.

5.3 Tit-for-Tat updation

For TFT, the technique used was that each individual adopted the strategy used by the majority of his neighbors on the previous turn. Since the changing of strategies was based purely on opponent strategies (and not on their payoffs), the values of P and R had no impact on the evolution

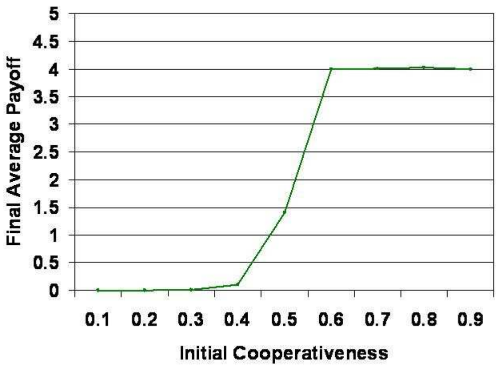

Figure 13: Stable Average Payoffs with TFT Updation

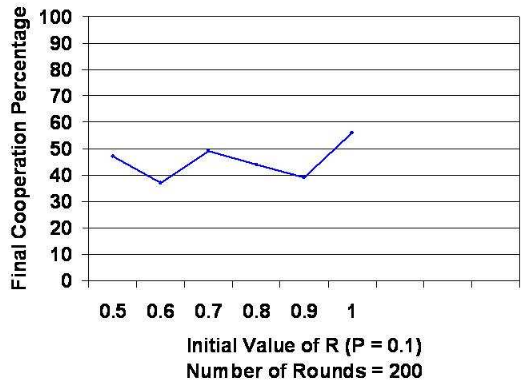

Figure 14: Approximate Stable Cooperation Levels with Pavlov Updation

of the population's strategy distribution. So, these were fixed at 0.4 and 0.6 respectively. Also, grid size was kept at 1000x1000, and number of rounds per simulation at 100. The initial cooperation percentage was varied from 0.1 to 0.9, and based on this, the convergent values of cooperation level and average payoff after 100 rounds were recorded. The results are shown in Figures 12 and 13.

As can be seen from Figure 12 , when cooperation level is below half to start with, it tends to disappear completely, whereas when it is initially above 0.5, we get a fairly high (though not necessarily 100 percent) cooperation level convergence. Also, the average payoff too nearly vanishes when cooperation level is below 0.5 initially, but converges to around 4.0 when cooperation is above 50 percent at the beginning. So, we see that the final situation depends entirely on how cooperative the population is to start with.

5.4 Pavlov updation

For Pavlov, also known as the win-stay, lose-shift strategy, the technique used was that any payoff of 4.0 or above was considered a win, and so strategy was switched only if a payoff of under 4.0 was obtained. This was in line with the technique used by Nowak and Sigmund (Nowak & Sigmund, 1993) for the two player Prisoner's Dilemma,where they showed that under certain conditions, Pavlov could beat TFT. Preliminary simulations showed that, like for Copycat, for Pavlov too the final state was not influenced by the initial cooperation level in the population. So the initial cooperativeness was fixed at 50 percent, while grid size was fixed at 500x500. Number of turns per simulation was kept

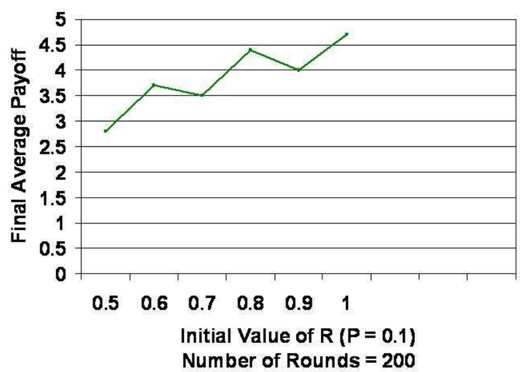

Figure 15: Approximate Stable Average Payoffs with Pavlov Updation

at 300. Simulations were run for different initial values of P and R. It was found that there was no clear dependence relation between these values and the final population state. Some amount of cooperation was always seen to emerge, but never a very high amount. The average payoff too tended to fluctuate anywhere in the 3.0-5.0 range. Results for one sample case, with initial value of P as 0.1 and initial value of R varying from 0.5 to 1.0, are shown in Figures 14 and 15.

It is evident from the results that the variance based on the initial P and R values is more or less random, though there is perhaps a slight increasing trend in both cooperativeness and average payoff, with increasing value of R(0). Increasing the initial value of P from 0.1 to 0.5 also showed no major change, only a similar slight increasing trend. Thus, these results suggest that if Pavlov is used, the final result is nearly the same irrespective of all the initial conditions, and cooperation does evolve to some extent, allowing for a decent average payoff to the population members.

5.5 Discussion of Results

We have seen how the cooperation level and average payoff evolves in the population under different kinds of learning or updation strategies being used by the members. From these results, it can be said that, in some sense, Pavlov is the 'best' updation model, as it seems to lead to approximately 50 percent cooperation and a decent average payoff, irrespective of the initial conditions. Of course, these values keep fluctuating, as members keep switching strategies whenever their payoff goes below 4.0, and this also means that this can be regarded as a 'fair' model, since in the long run everyone obtains similar overall benefit.

The copycat model, seemingly the most natural one in a situation such as this, was found to give good results in terms of cooperativeness and average payoff, only when the initial reward for mutual cooperation (i.e. the value of R) was above a certain threshold level. So, if there is insufficient incentive to cooperate at the start, the rate of defection goes down too fast, and cooperation is never able to evolve after that, even though the value of R subsequently increases by the evolution rule, as given in equation (10). The threshold levels could be changed by tweaking the values of k and $k^{'}$, to allow for faster or slower evolution of payoffs.

The TFT updation rule seems to be potentially the most damaging in this situation, as it will never allow cooperation to evolve in a population that does not initially have a cooperative majority. In cases where there is an initial cooperativeness of more than 50 percent, TFT does lead to a decent convergent average payoff of 4.0, but it may not always ensure fairness, as there may be individuals who are surrounded by defectors and thus cannot benefit as much as those having cooperative neighbors. In cases where the cooperation level is more than 0.5 to start with, but R(0) is below the copycat model threshold value for the given P(0), TFT will actually be better than copycat, but in general it is the least preferable of the three models.

Many changes are possible in the game model used here, and perhaps tweaking the values of parameters like k, $k^{'}$ and p_D^*, or using some more complex function for g(x), might give different results that may aid greater understanding of real-life social interaction. Also, this model does not take into account evolution of the population itself, and this is something else that may be looked at in future work. On the whole, the results show that all three considered updation strategies can lead to the evolution of cooperation, and thus any one of them, or quite likely some kind of mixture of all 3, may well be in use in human society in general.

6 Conclusions

We have looked at two games with dynamic payoffs, and tried to discover the equilibrium states reached in these games, and how the strategies and payoffs of the population members evolve with time. Both these games are close to very pertinent real-life situations - human language acquisition and human social interaction. The significance of our work lies in having shown that such situations can be modeled as dynamic-payoff games, and that the study of these games can lead to useful results pertaining to the given situations. Not much work has been done in relation to such games, and we hope that our very limited work can be extended in the future to take into account more of the variables that are implicit in the corresponding real-world 'games', in order to better understand them from a game-theoretic perspective.

7 Acknowledgments

We would like to thank Dr. Pabitra Mitra and Dr. Amitabha Mukerjee for their guidance during the course of this project. For fruitful discussions on game theory, we thank Dr. Pabitra Mitra, Dr. Mainak Chaudhuri, Utsav, Mathew and Maunendra.

References

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambridge, MA, USA: MIT Press.

Jain, S., Osherson, D., Royer, J. S., & Sharma, A. (1999). *Systems that learn.* Cambridge, MA, USA: The MIT Press. 2nd edition.

Komarova, N. L., Niyogi, P., & Nowak, M. A. (2001). The evolutionary dynamics of grammar acquisition. *Journal of Theoretical Biology*, *209*, 43–59.

Komarova, N. L., & Nowak, M. A. (2001). Natural selection of the critical period for language acquisition. *Proceedings of The Royal Society of London* (pp. 1189–1196).

Komorava, N. L., & Nowak, M. A. (2002). Population dynamics of grammar acquisition. In A. Cangelosi and D. Parisi (Eds.), *Simulating the evolution of language*. New York, NY, USA: Springer-Verlag New York.

Neumann, J. V., & Morgenstern, O. (1947). *Theory of games and economic behaviour*. Princeton, NJ, USA: Princeton University Press.

Niyogi, P. (2004). *The computational nature of language learning and evolution.* In press.

Nowak, M., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*, *364*, 56–58.

Osborne, M. J., & Rubenstein, A. (1994). *A course in game theory.* Cambridge, MA, USA: MIT Press.

Tomochi, M., & Kono, M. (2002). Spatial prisoner's dilemma games with dynamic payoff matrices. *Physical Review E*, *65*.