National University of Ireland, Galway
Ollscoil na hÉireann, Gaillimh

Department of Information Technology

technical report NUIG-IT-240900

The Voter's Paradox? Evolution of Cooperation in N-player games

S. Bradish (NUI, Galway)
C. O' Riordan (NUI, Galway)

The Voter's Paradox? Evolution of Cooperation in N-player games

S. Bradish & C. O' Riordan,
Department of Information Technology,
National University Of Ireland.

September 24, 2000

Abstract

Social dilemmas have been research by people in many fields. One common question or problem regards the existence of cooperation or altruism in situations where selfish behaviour appears to be of more benefit to the individual. Much attention has been payed to interactions involving 2 players where it has been shown that cooperation (altruism) can be beneficial in the long term. It is less obvious why cooperation exists in many-player interactions. This paper provides results to indicate that cooperation is the better strategy in environments where interactions vary from those involving 2 players to those involving many players. Given a sufficiently high proportion of 2-player games, cooperation will evolve as the norm. This paper describes initial results in this ongoing research project.

1 Introduction

This papers presents results regarding the problem of cooperation in multi-way interactions involving many players. Why does cooperation exist and how does it evolve, particularly in situations where it appears that selfish behaviour results in a higher gains for the involved individual? The problem of behaviour in many-many interactions is a pertinent one with many applications. The understanding of strategies in this setting may provide us with a better understanding of many such real-world settings and related phenomena— for example the tragedy of the commons, environmental issues and competition (and cooperation) between states and companies.

In this paper, we look at formal abstractions of multi-way interactions to attempt to show how cooperative behaviour can emerge. The majority of the research in the evolution of cooperation has dealt with cooperation in the two player case.

The paper is organised as follows: Section 2 deals with previous work in 2 player interactions (prisoner's dilemma). The third section deals with the evolution of strategies for these 2-player interactions .

The fourth section looks at typical interactions that involve more than 2 players, for example, the Voter's paradox[1]. These situations capture the essence of multi-way interactions. We illustrate the difficulty in evolving cooperation in such environments. Given the benefits to be gained by not cooperating, cooperation does not seem to be the best strategy. Yet in real world scenarios, we do see people behaving in an altruistic cooperative manner (people vote, people worry about and protect their environment, charity organisations exist).

Finally, we discuss some constraints that will lead to the evolution of cooperation as the norm.

[1] The Voter's paradox captures many of the salient features of a class of social dilemmas. Each individual that cooperates receives a reward. Each individual that does not cooperate also receives the same reward, but is spared any cost involved in cooperating

We illustrate that by imposing constraints on the frequency of occurrences of games (to mirror real world situations enforced by geographical limitations) cooperative behaviour can emerge as the social norm.

The final section summarises our results and discusses future work.

2 Prisoner's Dilemma and the Iterated Prisoner's Dilemma

The most famous dilemma that captures the essence of many 2-player interactions is the prisoner's dilemma.

In the prisoner's dilemma game, there are two players who are both faced with a decision—to either cooperate or defect. The decision is made by the player with no knowledge of the other players choice. If both cooperate, they receive a specific punishment or reward. If both defect they receive a larger punishment. However, if one defects, and one cooperates, the defector receives no punishment and the cooperator receives a punishment (the sucker's payoff). The game is often expressed in the canonical form in terms of pay-offs:

		Player 1	
		C	**D**
Player 2	**C**	(3,3)	(0,5)
	D	(5,0)	(1,1)

where the pairs of values represent the pay-offs for players `Player 1` and `Player 2` respectively. The prisoner's dilemma is a much studied problem due to it's far-reaching applicability in many domains. The prisoner's dilemma and applications has been described in [7][8][11] (biology) [13] (economics) and [4] (politics).

In the iterated version, 2 players will play numerous games (neither player knowing the exact number in advance). The probability of future interactions changes the nature of interactions between players. If there is a high probability of many future interactions, then cooperation (for fear of retaliation) would appear to be a more beneficial strategy. Some research has indicated that it is not necessary to look to the iterated versions for cooperation to occur. Work by Epstein[9] into spatial zones indicate that cooperative behaviour can emerge and exist in the non-iterated version of the game.

A computer tournament (Axelrod)[1] was organised to pit strategies against each other in a round-robin manner. Each strategy received a score for each run of the iterated prisoner's dilemma (IPD). The winning strategy was `tit-for-tat`; this strategy involved cooperating on first move and then mirroring opponents move on all subsequent moves. The initial results and analysis (which were echoed in later tournaments) showed that the following properties seemed necessary for success—niceness (cooperate first), retaliation, forgiveness and clarity. (The final property has been questioned in [3], who develop a strategy that is nice, forgiving and retaliatory but not clear).

In a second tournament[1], of the top 16 strategies, 15 were found to be nice. These results seem to indicate that cooperative strategies will flourish if there is a high chance the strategies will meet again.

3 Evolution of Strategies for the Iterated Prisoner's Dilemma

Axelrod ran artificial life simulations by pitting strategies against each other. The fitness of a strategy was measured by its score in interactions; a strategy's representation in the next generation was proportional to this fitness. No best strategy exists; the success of a strategy depends on the other strategies present in the population.

Results showed that lower ranked strategies (which were predominately uncooperative) died off quickly, where strategies continued to flourish. The only highly ranked non-nice strategy survived temporarily but then quickly died off as the naive strategies upon which it preyed, became extinct.

Other experiments have also been carried out where the initial strategies are created randomly. A genetic algorithm[10] is then used to breed newer, more successful strategies (again, cumulative score in the IPD is taken as a measure of fitness).

Different approaches have been taken to represent the strategies including:

- (i, p, q) framework described in Nowak and May[12] and in Cohen and Axelrod [6]. Three values are maintained: i, the probability of cooperating on the first move, p, the probability of cooperating following an opponent's cooperation and q, the probability of defecting following a defection by an opponent. For example, $i = p = 1, q = 0$ would express the `tit-for-tat` strategy.
- Moore's machine (1, 2 game histories). A binary string is used to represent the strategies' action based on a game history. For example, in a one-game history 3 bits are required to represent the opening move, the move to play following a defection in the last game and the move to play following a cooperation in the last game. In a 2-player game, in addition to opening moves, moves to play following the combinations (CC, CD, DC, DD) are required.

Many such experiments arrive at similar results and conclusions—strategies like `tit-for-tat` evolve.

4 Multi-way Interactions

Many-player games involve a set of players making a decision (to cooperate or defect). A pay-off matrix exists to specify the rewards for each of the players given the number of cooperators and defectors.

One prime example is the Voter's paradox where a benefit is returned to all members in the group, irrespective if the member contributes or not. Given this scenario, it is not unreasonable to expect non-cooperation to be the favoured strategy, yet there are many examples in nature and society where cooperation is the favoured strategy. Elections, environmental actions and the tragedy of the commons are all examples of this phenomenon.

We attempt to model the above scenario as a n-player matrix: all get a payoff if all contribute, all contributors pay a certain cost; defectors pay nothing.

In this case, defective behaviour is favoured and there is little pressure for cooperative behaviour. If we apply evolutionary computing approaches to developing successful strategies, defective strategies will evolve. It has been argued that defective behaviour emerges in real-world scenarios due to the obvious benefits to be gained by defection. Furthermore, in many-player games anonymity is often guaranteed, so there is no fear of reprisal in future interactions. It is more difficult to explain why we encounter cooperation in the real-world examples of these social dilemmas.

A few attempts have been made to attempt to explain, via computer simulations why cooperation can emerge and exist.

Simulations by Axelrod[2] illustrate that norms of cooperation can be created and enforced. Axelrod created a model that allows players in the model to punish those that break the norm of cooperation. This punishment allows the emergence of cooperation but does not guarantee the stability of cooperation as a norm. A further extension, where those agents who do no actively punish defections can also be punished allows the emergence of cooperation and allows the stability of cooperation as a norm.

In our experiments, we attempt to model many player games (ranging from 2 player games to games involving all N players). We hypothesise that cooperative behaviour will emerge as the favoured strategy. We create pay-off matrices for each of the N-1 game-types. We show that given certain frequencies of occurrences of games, cooperative strategies will be established as the norm. We adopt a genetic algorithm to evolve strategies.

5 Experimental Setup

5.1 Extending the range of games

Games may involve interactions between 2 players (the classical prisoner's dilemma and the iterated prisoner's dilemma) or interactions between many players (Voter's paradox). We wish to allow the players in our game to partake in a range of games: from 2 player to N-player. To model real world scenarios we allow 2 player games to occur more frequently than 3-player etc. This seems to be a valid assumption if one considers the effect geographical locations play in real world scenarios.

More formally, for any, i, the frequency of games involving i players is greater than the frequency of games involving $i+1$ players.

We chose the following implementation: given n players, we have $n-1$ game types. We allow the 2 player games to account for 50% of the games, 3 player games to account for 25% of the games etc. This can be easily implemented by choosing a random number in the range 1—2^{n-1} (akin to roulette wheel selection commonly used in genetic algorithms.)

For each of these games, we require a corresponding pay-off matrix to determine rewards for cooperators and defectors. In a two-player game it should be evident that cooperation will be rewarded. In many player games, like the Voter's paradox, if there is a reward (enough voters) there will be an equal payoff for all creatures, cooperators and defectors alike. Cooperators of course have the extra cost of cooperating.

5.2 Pay-off Matrices

We formulate the above requirements as follows:

We let the pay off for players be divided equally between the players. Hence:

$$\text{pay-off} = \frac{R}{N}$$

where R is total reward and N is the number of players.

The more player contributing (cooperating), the greater the reward. Therefore we set R to be a function of the number of cooperators:

$$R = f(|C|)$$

A cost is involved for all players. For defectors, the cost is 0, for cooperators the cost represents the effort involved in cooperating. We let the cost be greater if fewer cooperate (this is realistic for some of the well-known social dilemmas). Therefore the more defectors there are, the greater the cost for the cooperators:

$$\text{cost} = \begin{cases} 0, \text{for defectors} \\ \frac{D}{N}, \text{for cooperators} \end{cases}$$

$R = f(|C|)$, R is a function of the number of cooperators.

$cost = 0$, for defectors

The following will satisfy the above constraints

For cooperators:

$$\text{pay-off} = (\frac{C}{N} \times N^2) - (\frac{D}{N} \times N^2)$$

For defectors:

$$\text{pay-off} = \frac{C}{N} \times N^2$$

It is evident that in a 2-player games, cooperation is favoured, in 3-player games and higher, defection becomes the favoured approach.

5.3 Representation of strategies

We represent strategies using a 6 bit representation. The first bit represents the first move to be played (cooperate or defect). The other 5 bits represent the strategies response to the following cases:

- all cooperated in the last game
- all defected in the last game
- equal number of defectors/cooperators in the last game
- more defectors than cooperators
- more cooperators than defectors

We employ a genetic algorithm (one-point crossover, 0.01 mutation, initial population size = 100) to evolve strategies. Rewards in the game are used as a measure of fitness.

6 Results

Initially we didn't enforce any constraints on the frequency of the games. As expected, defection became the norm quite quickly.

Upon imposing constraints of the frequencies of the games, we see that cooperators tend to flourish. Figure 1 plots the number of strategies that cooperate.

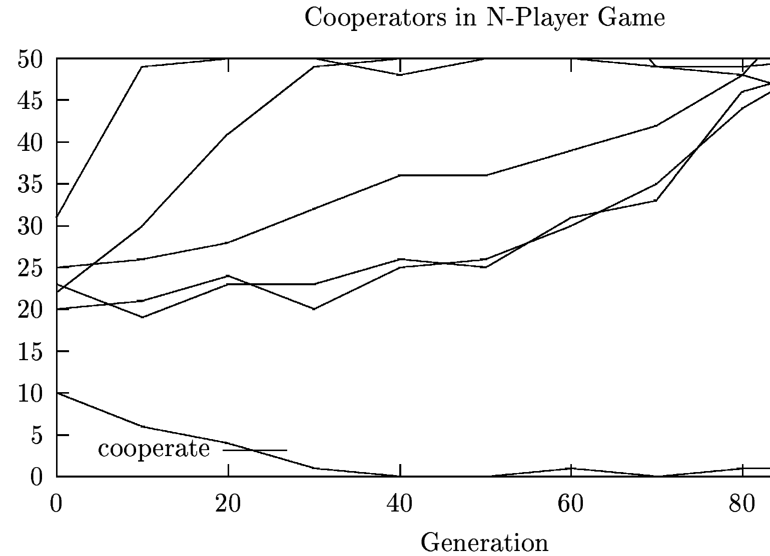

A more accurate measure of the number of cooperators is to count/examine the number of strategies that follow a cooperation by an opponent with a cooperation. Upon examination we see a large number of `tit-for-tat` like strategies.

We also see a number of spiteful strategies—those that will cooperate initially, return cooperation with cooperation provided no defection occurs. Upon defection, these strategies always defect.

7 Future Work

One aspect of the above work that could be further investigated is the effect of different pay-off matrices. The formalisms we use represents just one possible approach to capturing the essence of common social dilemmas. Other formalisms have been used to describe pay-off matrices for N-player games most notably[5]. Future work may involve experimenting with these formulae. At present the above formalism sufficiently captures the nature of N-player dilemmas.

One criticism that could be leveled at our model, is the strict enforcement of the frequency of games (half the games are 2-player games). We are currently running further simulations to see when cooperation breaks down when the frequency of games involving many players are increased.

Another experiment we are currently investigating is to augment the strategy representation in our genetic algorithm to include a set of bits that represent a 'game-type' which would represent a strategies' willingness to take part in certain types of games. We hypothesise that games involving few players will emerge as the favoured games providing evidence to support the prevalence of small groups in social organisations.

8 Conclusion

Despite the advantages to be gained and the pressure to adapt non-cooperation as the norm in N-player games, we have shown without building punishment mechanisms into our model that cooperation can emerge.

The only enforced constraint in our model is that of frequency of types of games. We allow 50% of the games to be 2-player games (in which cooperation is favoured) the rest involving more than 2 players (in which defection is favoured).

References

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

[2] Robert Axelrod. *The Complexity of Cooperation : Agent-Based Models of Competition and Collaboration.* Princeton Univ Pr, 1997.

[3] B. Beaufils, J. P. Delahaye, and P. Mathieu. Our meeting with gradual, A good strategy for the iterated prisoner's dilemma. In C. Langton, editor, *Artificial Life, Proc. 5th International Workshop on the Synthesis and Simulation of Living Systems*, volume 5, 1996. Artificial Life 5, Nara, Japan, May 16-18.

[4] Jonathon Bendor and Dilip Mookherjee. Institutional structure and the logic of ongoing collective action. *American Political Science Review*, 81:129–154, 1987.

[5] R. Boyd and P. Richerson. The evolution of reciprocity in sizable groups. *Journal of Theoretical Biology*, 132:337–356, 1989.

[6] Micheal Cohen, Rick Riolo, and Robert Axelrod. The emergence of social organization in the prisoners' dilemma: How context-preservation and other factors promote cooperation. Technical Report PSCS-98-12-001, Center for the Study of Complex Systems (CSCS), University of Michigan, 1998.

[7] L.A. Dugatkin. Do guppies play tit for tat during predator inspection visits? *Behavioral Ecology and Sociobiology*, 23:395–399, 1988.

[8] L.A. Dugatkin. N-person games and the evolution of co-operation: A model based on predator inspection in fish. *Journal of Theoretical Biology 142:123-135*, 142:123–135, 1990.

[9] Joshua M. Epstein. Zones of cooperation in demographic prisoner's dilemma. Technical report, Santa Fe Institute, 1997.

[10] John Holland. *Adaptation in natural and artificial systems.* Ann Arbor, 1975.

[11] Michael Lombardo. Tree swallows and tit for tat: Response to koenig. *Ethology and Sociobiology*, 11:521–528, 1990.

[12] M.A. Nowak and R.M. May. Evolutionary games and spatial chaos. *Nature*, 335(827), 1992.

[13] Elliot Sober. Stable cooperation in iterated prisoner's dilemmas. *Economics and Philosophy*, 8:127–139, 1992.