National University of Ireland, Galway
Ollscoil na hÉireann, Gaillimh

Department of Information Technology

technical report NUIG-IT-260601

Iterated Prisoner's Dilemma: A review

C. O'Riordan (NUI, Galway)

Iterated Prisoner's Dilemma: A review

Colm O' Riordan
Department of Information Technology
National University of Ireland, Galway
Ireland.
colmor@geminga.nuigalway.ie

Abstract

Much debate has centered on the nature of social dilemmas. In environmental issues, trade-wars, negotiations between countries and social interactions, there are many scenarios where involved players may choose altruistic, cooperative action or an alternative selfish behaviour.

Games, particularly the prisoner's dilemma and variations have been used to model and analyse such scenarios.

In this paper, an overview of research in this domain is presented. The iterated prisoner's dilemma (IPD) is discussed with attention payed to evolution of strategies in this environment, the effect of noise in the IPD and finally N-player versions.

1 Introduction

Much debate has centered on the nature of social dilemmas. In environmental issues, trade-wars, negotiations between countries and social interactions, there are many scenarios where involved players may choose altruistic, cooperative action or an alternative selfish behaviour. In many cases, it appears that the selfish action is more beneficial and is the rational choice to take; however we encounter time and time again, people (and groups) choosing the apparently non-rational altruistic option. Examples include people investing time, money and effort in environmental issues and charity organisations.

Thomas Hobbes[18] presented a rather pessimistic explanation as to how cooperation can be maintained in a group of interacting agents; he argued that prior to the existence of governments that nature was dominated by selfish agents resulting in a life that was "solitary, poor, nasty, brutish and short". He believed that without a controlling authority, cooperation was impossible. Olson[35] agreed with the classical conclusion that coercion or selective incentives are necessary to achieve cooperation.

More recent work, most notably that of Robert Axelrod[1], has shown that cooperation can emerge as a norm in a society comprising entities with individual selfish motives.

The most oft studied games in this domain are the Prisoner's Dilemma (PD) and variations such as the Iterated Prisoner's Dilemma (IPD). The prisoner's dilemma captures, in an abstract manner, the salient features of many social dilemmas rendering it a fascinating problem to analyse in attempting to understand phenomena such as cooperation, altruism and free-riding. In this paper, research and results in the prisoner's dilemma are reviewed.

2 Prisoner's Dilemma

In the prisoner's dilemma game, two players are both faced with a decision—to either cooperate(C) or defect(D). The decision is made by a player with no knowledge of the other player's choice. If both cooperate, they receive a specific punishment. If both defect they receive a larger punishment. However, if one defects, and one cooperates, the defecting strategy receives no punishment and the cooperator 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**	(λ_1, λ_1)	(λ_2, λ_3)
	D	(λ_3, λ_1)	(λ_4, λ_4)

where the pairs of values represent the pay-offs (rewards) 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. In game theory, the prisoner's dilemma can be viewed as a two-person, non-zero-sum, non-cooperative and simultaneous game. In order to have a dilemma the following must hold: $\lambda_3 < \lambda_1 < \lambda_4 < \lambda_2$, where λ_2 is the sucker's payoff, λ_4 is the punishment for mutual defection, λ_1 is the reward for mutual cooperation and λ_3 is the temptation to defect. The constraint $2\lambda_1 > \lambda_2 + \lambda_3$ also holds.

The prisoner's dilemma and applications has been described in many domains including biology[12][13][25], economics[38] and politics[6].

3 Iterated Prisoner's Dilemma

The game becomes more interesting, and more widely studied, in the iterated version where 2 players will play numerous games (the exact number not known to either player). Each player adopts a strategy to determine whether to cooperate or defect at each of the moves in the iterated game.

3.1 Strategies

Before discussing the main results obtained in the iterated prisoner's dilemma it may be instructive to try to classify the strategies.

periodic: strategies play C or D in a periodic manner. Common strategies: *ALL-C*, *ALL-D*, *(CD)**, *(DC)**, *(CCD)**, etc.

random: strategies that have some random behaviour. Totally random, or one of the other types (e.g. periodic) with a degree of randomness.

based on some history of moves: *tit-for-tat* (C initially, then D if opponent defects, C if opponent cooperates), *spiteful* (C initially, C as long as opponent cooperates, then D forever), probers (play some fixed string, example (DDC) and then decides to play *tit-for-tat* or *ALL-D* (to exploit non-retaliatory), *soft-majo* (C initially, then cooperate if opponent is not defecting more than cooperating).

There are many variations on each of the above type of strategies.

3.2 Results

A computer tournament[1] was organised to pit strategies against each other in a round-robin manner in an attempt to identify successful strategies and their properties. The winning strategy was *tit-for-tat* (TFT); this strategy involved cooperating on first move and then mirroring opponents move on all subsequent moves.

The initial results and analysis showed that the following properties seemed necessary for success—niceness (cooperate first), retaliatory, forgiving and 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 are useful if there is a high chance the strategies will meet again.

Further analysis involved the development of a genetic algorithm to evolve successful strategies. The more successful strategies tended to be more complex than the traditional TFT and violated the fourth heuristic (that of clarity) proposed by Axelrod :"Don't be too clever"; these strategies are quite complex.

Beaufils et al[3] question that last property and develop a strategy *gradual*[1] which is far more complex than *tit-for-tat* and outperforms *tit-for-tat* in experiments.

The *forgiving*[33] strategy also challenges the final property; *forgiving* is not clear or simple and has proven strong in environments similar to those used by Beaufils[3].

No best strategy exists; the success of a strategy depends on the other strategies present. For example, in a collection of strategies who defect continually (*ALL-D*) the best strategy to adopt is *ALL-D*. In a collection of strategies adopting a *tit-for-tat* strategy, an *ALL-D* strategy would not perform well.

3.3 Stability

This type of reasoning leads to the question of which strategies can continue to exist—do stable strategies exist? This question was first addressed by Maynard-Smith[37] who introduced the concept of an evolutionary stable strategy—a strategy is (collectively) stable if no other strategy can *invade* it; by *invade* we mean that a new strategy scores higher interaction with a native strategy than a native strategy does interacting with a native strategy.

An evolutionary stable strategy must be a 'best response' to itself, because a mutant playing a better response would

[1] gradual performs like *tit-for-tat*, in that it cooperates on the first move. It retaliates upon defection. On the first defection it responds with a defection(D), followed by 2 cooperations(CC). Following the second defection, it responds with 2Ds, followed by 2Cs and so forth.

have a higher average payoff[37].

Maynard-Smith[37] defines the concept of a strongly stable strategy as follows: A strategy is strongly stable if it will drive any invaders to extinction. More formally, a strategy i is strongly stable if:

$$V(i,i) > V(j,i) \text{ or}$$
$$V(i,i) = V(j,j) \text{ and } V(i,j) > V(j,j)$$

where i and j are strategies and $V(i,j)$ is the reward obtained by i when playing against strategy j.

For any strategy i, there can exist strategies that are different from i but perform against strategy i and itself as strategy i would. Hence, Maynard-Smith's definition of strong stability can not hold.

A strategy is weakly stable if it will not be driven to extinction by any invaders. A strategy i is weakly stable if:

$$V(i,i) > V(j,i) \text{ or}$$
$$((V(i,i) = V(j,j) \text{ and } V(i,j) \leq V(j,j))$$

This definition was first proposed by Bendor and Swistak[7].
The evolution of cooperative can be viewed as three separate questions[2]:

- Robustness: what type of strategy can thrive in a variegated environment?
- Stability: under what circumstances can such a strategy, once fully established, resist invasion by mutant strategies?
- Initial viability: even if a strategy is robust and stable, how can it get a foothold in an environment which is predominantly non-cooperative?

Attempts have been made to answer the above questions via simulations. The final question requires simulations that mimic "survival of the fittest" behaviour. These experiments which breed strategies in an evolutionary setting are reviewed in the next section.

4 Evolution of strategies

4.1 Introduction

A natural means to explore more fully the range of strategies and their performance is to adopt ideas from the domain of evolutionary computation. In these approaches, one attempts to evolve suitable strategies. The most well-known approach is that of genetic algorithms[21], where a population of solutions is created and then subjected to the process of evolution. The initial population (randomly created) represents a set of strategies. The process mimics Darwin's theory of evolution ('survival of the fittest'), where fitter solutions are passed on the subsequent generations possibly subjected to the operations of mutation and crossover (mating) to create newer, possibly fitter, solutions.

The question of representation of strategies is an important one—how should strategies for the iterated prisoner's dilemma be represented in a way suitable for genetic algorithms and other evolutionary algorithms .

We outline, briefly, some approaches and results:

Nowak and Sigmund[32] use a triple to represent a set of reactive strategies. The triple of values indicate, respectively, the probability of cooperating on the first move, the probability of cooperating following an opponent's cooperation and the probability of cooperating following a defection. `tit-for-tat` can be represented as $(1, 1, 0)$.

In most experiments, defection evolved as the norm. When one of the initial strategies are close to tit-for-tat, cooperation can flourish, but having eliminated the exploiters, tit-for-tat is superseded by a strategy closer to GTFT[2]

A more expressive representation used in later experiments involved the use of four probabilities indicating the probability of a strategy cooperating following a (C,C), (C,D), (D,C) and (D,D) respectively. Following simulations, states of mutual cooperation were evolved. Over 90% of strategies were of the type Pavlov[23][3]. The remaining comprised tit-for-tat and GTFT.

Linster[24] uses Moore machines to represent strategies. In these experiments it was shown that no one strategy dominated the environment. The most successful strategy was spiteful[4] (aka GRIM, TRIGGER).

Beaufils[4] attempt to analyse the different classes of strategies for the iterated prisoner's dilemma using evolutionary computation techniques. Three different types of genotypes (which encode the phenotype) are explored:

[2] GTFT, initially described by Molander[29] has probability of cooperating following a cooperation equal to 1, and probability of cooperating following a defection equal to $\min((1 - \frac{(\lambda_3 - \lambda_1)}{(\lambda_1 - \lambda_2)}), \frac{(\lambda_1 - \lambda_4)}{(\lambda_3 - \lambda_4)})$

[3] The Pavlov strategy acts as follows: if on the previous move both strategies performed the same action, Pavlov cooperates; otherwise it defects

[4] Spiteful cooperates on its first move and continues to cooperate as long as its opponent does so; following a defection, spiteful continues to defect)

- memory: each strategy can see past moves.
- binary memory: similar to the previous, but in this case the move is based on past moves and whether the opponent has defected more often than it has co-operated.
- memory automata: represents a two state automata.

This approach presents a mechanism to represent a large set of strategies in an unbiased means.

Each genotype contained 19 genes representing various features of the strategy—the first move to make, random defection, how to detect defection, punishment thresholds, forgiveness, etc. Some interesting results, other than the evolution of complex successful strategies were also found, namely:

- the gene determining the first move converged to 1 indicating that the strategy should always cooperate on the first move.
- the gene determining the degree of random defection quickly converged to 0 indicating that unprovoked defection is a bad idea.
- forgiveness is not always a good idea; strategies should not forgive defecting strategies.

Harrald and Fogel[16] use single-layer feed-forward neural networks to represent strategies. The inputs (six in total) to the neural network represent the previous three moves by both players. All networks played against each other in a round-robin manner; the fitness assigned to the networks is a function of the score obtained by the network. The networks in the fitter half of the population were chosen to become parents in the next round.

Cohen et al.[10] report on the factors that promote cooperation. They attempt to explore the emergence and maintenance of cooperation by analysing the effect of 3 factors: strategy space, interaction processes and adaptive processes. They show that cooperation can emerge given random mixing of agents. The three dimensions are represented as:

- Strategy Space
  1. Binary strategy: each strategy is represented as a triple (i, p, q), i representing the first move, p representing the probability that the strategy co-operates, q the probability that the strategy defects (p and $q \in \{0, 1\}$)
  2. Continuous: as above but p and q may range across the interval $[0, 1]$
- Interaction Processes: a set of 6 interaction types are defined—including fixed neighbourhoods, randomly selected etc.
- Adaptive Processes: specifies techniques by which a strategy can adapt over time—by imitation, by using best solution encountered so far incorporated via GA operators

5 Noisy environments

The majority of work in the iterated prisoner's dilemma has focussed on the games in a noise-free environment, i.e., there is no danger of a signal being misinterpreted by the opponent or the message being damaged in transit.

This assumption of a noise-free environment is not necessarily valid if one is trying to model real-world scenarios.

There are different means that can be chosen to introduce noise to the simulation:

- mis-implementation (when the player makes a mistake implementing its choice)
- mis-perception (when one player misperceives the other player's signal or choice)

Bendor[5] effects noise by introducing payoffs that are subjected to error. Upon cooperation in face of defection by an opponent, a person receives the payoff $S + e$, where e is random with expected value 0.

In [15], it is argued that "if mistakes are possible evolution may tend to weed out strategies that impose drastic penalties for deviations".

Kahn and Murnighan [22] find that in experiments dealing with prisoner's dilemma in noisy environments, cooperation is more likely when players are sure of each other's payoffs. Miller's experiments in genetic algorithms applied to the prisoner's dilemma results in the conclusion that cooperation is at its greatest when there is no noise in the system and that this cooperation decreases as the noise increases[28].

Some ideas to promote cooperation in environments have been posited by Axelrod; these include genetic kinship, clustering of like strategies, recognition, maintaining closeness when recognition capabilities are limited or absent (e.g limpets in nature), increasing the chance of future interactions (certain social organisations, hierarchies in companies etc.), changing the pay-offs, creating social norms where one learns cooperation.

Hoffman[20] reports that results are sensitive to the extent to which players make mistakes either in the execution of their own strategy (mis-implementation noise) or in the perception of opponent choices (mis-perception noise). In particular, cooperation is vulnerable to noise as it is supported by conditional strategies. For example, in a game between two TFTs, a single error would trigger a series of alternating defection. Axelrod (1984) repeated his initial round-robin tournament with added 1% chance of players misunderstanding their opponent's move in any round. He found that TFT still came first despite some echoes of retaliation between cooperative strategies.

It can be shown that higher degrees of noise can be detrimental to TFTs performance. Given noise of p percent, TFT against itself can be captured via the following transition matrix.

$$\begin{pmatrix} (1-p)^2 & 2(1-p)p & p^2 \\ (1-p)p & (1-p)^2+p^2 & (1-p)p \\ p^2 & 2(1-p)p & (1-p)^2 \end{pmatrix}$$

Solving above equations, we find that for any noise ($p > 0$), we get $\alpha = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$.

A number of authors confirm the negative effect of noise of TFT and find that more forgiveness promotes cooperation in noisy environments[5] Mueller [31].

Other interesting results are also reported. These include 'pavlovian' strategies which are more likely to avoid spirals of defection that `tit-for-tat`[23] (also shown to perform well in[27]), the lowering of levels of cooperation in a society without the introduction of defecting strategies[28][26], the effect of highlighting differences between strategies that would coexist in noise-free environments[8].

6 Spatial

Note that some research has indicated that it is not necessary to look to the iterated versions for more interesting behaviour to occur. Work by Epstein [14] into spatial zones indicate that more interesting behaviour (e.g mutual cooperation) can emerge and exist in the non-iterated version of the game.

Similar work on the effect of spatial organisation of strategies was undertaken by Oliphant[34] who showed, via a series of simulations, that in the absence of spatial constraints, the population quickly fell into defection. However with spatially constrained populations, it was possible to evolve and maintain cooperative behaviour.

7 Related Dilemmas and problems

7.1 Lift dilemma

Another game, closely related to the prisoner's dilemma, is the lift dilemma in which two players interact. The first equality associated with the prisoner's dilemma holds but the second inequality is changed. This leads to the scenario where there now exist two forms of cooperation—mutual cooperation as in the prisoner's dilemma and a more successful form where two strategies alternate between pairs of (C,D) and (D, C); in effect they take turns at obtaining the the sucker's payoff, in order to obtain the maximum pay-off on the next turn.

Delahaye and Beaufils[11] present REASON which plays a C initially with some probability on the first move and while the previous round is phased. On subsequent moves, this the strategy plays either $(C,D)^*$ or $(D,C)^*$. A variation on this REASON-TFT which plays C with some probability on the first move and while the previous round is phased; else tit-for-tat is played. Both of these flourish in an evolutionary simulation and can illicit the higher form of cooperation than that obtained with successful strategies in the traditional iterated prisoner's dilemma.

7.2 N-player dilemmas

One prime example of an N-player dilemma is the Voter's paradox "where it is true that a particular endeavour would return a benefit to all members where each individual would receive rewards, it is also true that any member would receive an even greater reward by contributing nothing". Elections, environmental actions and the tragedy of the commons are all examples of this phenomenon.

In such games, each player can choose to defect or cooperate and there is no external central control; in this scenario cooperation is costly and defection is cost-less, hence rational self-interested individuals should always defect, even if the group outcome from joint defection is not Pareto optimal. Common mechanisms to ensure cooperation is through reputation based schemes; however in many real-world scenarios players are largely anonymous. In these scenarios it is expected that the dominant strategy be one of defection.

In an n-player dilemma, each player faces a choice between two alternatives: to cooperate or to defect. The payoffs are functions of the number of cooperators. Let $V(c|i)$ denote the pay-off for cooperators and $V(d|i)$ the pay-off for defectors given i cooperators.

According to Molander[30] the following conditions should hold:

- monotonicity: $V(c|i) > V(c|i-1)$ and $V(d|i) > V(d|i-1)$, $i = 1, 2, \ldots n$
- dominance of the D alternative: $V(d|i) > V(c|i)$, $i = 0, 1, \ldots n-1$
- efficiency of cooperation $(i+1)V(c|i) + (n-i-1)V(d|i+1) > iV(c|i-1) + (n-i)V(d|i)$, $i = 1, 2, \ldots n$ and $V(c|n-1) > V(d|0)$

Boyd and Richerson [9] also tackle the problem of evolution in N-player games; they show that cooperation is more difficult to illicit with large groups. In their model, groups are formed by sampling N individuals from the population who interact in the a repeated n-person dilemma. Using a similar formalism to that above, from their analysis of strategies TFT and ALL-D, the authors conclude that cooperation is only ever to emerge in extremely small groups.

Recent work by O'Riordan and Bradish[36], simulates an environment where players are engaged in many different types of games ranging from the traditional 2-player game to games involving many players. Preliminary results show that cooperation can emerge given a high percentage of 2-player games.

8 Summary

This paper gives a brief overview of current and past research in the domain of the prisoner's dilemma. Included in the paper is a discussion of the prisoner's dilemma and the iterated version and in related versions (noisy environments and spatial constraints).

9 Acknowledgement

The author would like to express gratitude to Josephine Griffith for feedback and comments on drafts of this paper.

References

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

[2] R. Axelrod and W. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390–1396, March 1981.

[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] B. Beaufils, P. Mathieu, and J.-P. Delahaye. Complete classes of strategies for the classical iterated prisoner's dilemma. In *Proceedings of Evolutionnary Programming VII, LNCS vol 1447*, pages 33–41. Springer-Verlag, 1998.

[5] J. Bendor, R.M. Kramer, and S. Stout. When in doubt cooperation in a noisy prisoner's dilemma. *Journal of Conflict Resolution*, 35(4):691–719, 1991.

[6] J. Bendor and D. Mookherjee. Institutional structure and the logic of ongoing collective action. *American Political Science Review*, (81):129–154, 1987.

[7] J. Bendor and K.P. Swista. The evolutionary stability of cooperations. *American Political Science Review*, 91(2):290–307, june 1997.

[8] R. Boyd. Mistakes allow evolutionary stability in the repeated prisoner's dilemma. *Journal of Theoretical Biology*, 136:47–56, 1989.

[9] Robert Boyd and Peter Richerson. The evolution of reciprocity in sizable groups. *Journal of Theoretical Biology*, 132:337–356, 1988.

[10] Michael D. Cohen, Rick L. Riolo, and Robert Axelrod. The emergence of social organization in the prisoner's dilemma: How context-preservation and other factors promote cooperation. Technical report, SantaFe Institute, 1999.

[11] J. P. Delahaye and P. Mathieu. Random strategies in a two levels iterated prisoner's dilemma : How to avoid conflicts ? In H. J. Mller and R. Dieng, editors, *Proceedings of the ECAI 96 Workshop: Modelling Conflicts in AI*, pages pages 68–72, 1996.

[12] L.A. Dugatkin. Do guppies play tit for tat during predator inspection visits? *Behavioral Ecology and Sociobiology*, 2(3):395–399, 1988.

[13] 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, 1990.

[14] J. Epstein. Zones of cooperation in demogra phic prisoner's dilemma. Technical report, Santa Fe Institute, 1997.

[15] Drew Fundunberg and Eric Maskin. Evolution and cooperation in noisy repeated games. *The American Economic Review*, 80(2):274–179, 1990.

[16] Paul Harrald and David Fogel. Evolving continuous behaviours in the iterated prisoner's dilemma. *Biosystems*, 37:135–145, 1996.

[17] J. Hirshleifer and J.M. Coll. What strategies can support the evolutionary emergence of cooperation? *Journal of Conflict Resolution*, 32(2):367–398, 1988.

[18] Thomas Hobbes. *Leviathan*. 1651.

[19] R. Hoffman and N. Waring. *Complexity Cost and Two Types of Noise in the Repeated Prisoner's Dilemma*. Springer Verlag, 1998.

[20] Robert Hoffman. Twenty years on: The evolution of cooperation revisited. *Journal of Artificial Societies and Simulation*, 3(2), 2000.

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

[22] L.M. Kahn and J.K Murnighan. Conjecture, uncertainty, and cooperation in prisoner's dilemma games: Some experimental evidence. *Journal of Economic Behaviour & Organisation*, 22:91–117, 1993.

[23] D. Kraines and V. Kraines. Learning to cooperate with pavlov an adaptive strategy for the iterated prisoner's dilemma with noise. *Theory and Decision*, 35:107–150, 1993.

[24] B.G. Linster. Evolutionary stability in the infinitely repeated prisoner's dilemma played by two-state moore machines. *Southern Economic Journal*, (58):880–903, 1992.

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

[26] B. Lomborg. Nucleus and shield: The evolution of social structure in the interated prisoner's dilemma. *American Sociological Review*, (61):278–307, 1996.

[27] L. Marinoff. Maximizing expected utilities in the prisoner's dilemma. *Journal of Conflict Resolution*, 36(1):183–216, 1992.

[28] J.H. Miller. The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behaviour & Organisation*, 29:87–112, 1996.

[29] Per Molander. The optimal level of generosity in a selfish, uncertain environment. *Journal of Conflict Resolution*, 29:611–619, December 1985.

[30] Per Molander. Tthe prevalance of free riding. *Journal of Conflict Resolution*, 36(4):756–771, December 1992.

[31] U. Mueller. Optimal retaliation for optimal cooperation. *Journal of Conflict Resolution*, 31(4):692–724, 1988.

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

[33] Colm O' Riordan. A *forgiving* strategy for the iterated prisoner's dilemma. *Journal of Artificial Societies and Social Simulation*, 3(4), 2000.

[34] M. Oliphant. Evolving cooperation in the non-iterated prisoner's dilemma. *Artificial Life IV*, 1994.

[35] M. Olson. *The logic of collective action: public goods and the theory of groups*. Harvard University Press, Cambridge, MA, 1965.

[36] C. O' Riordan and S. Bradish. Experiments in the iterated prisoner's dilemma and voter's paradox. *11th Irish Conference on Artificial Intelligence and Cognitive Science*, 2000.

[37] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.

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