

Our Meeting With Gradual: A *Good* Strategy For The Iterated Prisoner's Dilemma

Bruno Beaufils, Jean-Paul Delahaye and Philippe Mathieu

Université des Sciences et Technologies de Lille

Laboratoire d'Informatique Fondamentale de Lille – U.R.A. 369 C.N.R.S.

U.F.R. d'I.E.E.A. Bât. M3 – F-59655 Villeneuve d'Ascq Cedex

e-mail: {beaufils, mathieu, delahaye}@lifl.fr

Abstract

In this paper, after a short return to the description of the classical version of the Iterated Prisoner's Dilemma and its application to the study of cooperation, we present a new strategy we have found named *gradual*, which outperforms the *tit-for-tat* strategy, on which are based a lot of works in the Theory of Cooperation. Since no pure strategy is evolutionarily stable in the IPD, we cannot give a mathematical proof of the absolute superiority of *gradual*, but we present a convergent set of facts that must be viewed as strong experimental evidences of the superiority of *gradual* over *tit-for-tat* in almost every rich environment. We study in detail why this strategy is a good one and we try to identify the difference it has with *tit-for-tat*. Then we show how we have improved the strength of *gradual* by using a genetic algorithm, with a genotype we have created, which includes a lot of well-known strategies for the IPD, such as *tit-for-tat*. We present our ideas on a tree representation of the strategies space and finally we propose a new view of evolution of cooperation in which complexity plays a major role.

1 The Classical Iterated Prisoner's Dilemma

In this paper we first present the classical version of the Iterated Prisoner's Dilemma which is fully described in [30] and its application to the study of cooperation. We insist on the importance of the pure version of this problem. Then we present a new strategy we have found called *gradual*, which outperforms the *tit-for-tat* strategy which is classically recognized as the best possible strategy in the classical IPD.

In section 2, we present in details the *gradual* strategy which has the same qualities than *tit-for-tat* plus a progressive retaliation and forgiveness possibilities. Two experiments are reported, the first one with twelve general strategies to show the precise behavior of it, the second one with the results of a tournament we have organized in the French edition of the "Scientific American".

In section 3 we try to improve the strength of this strategy by using a genetic algorithm, on a genotype we have created and which includes lots of well-known strategies (in fact our genotype can cover more than 8×10^{15} strategies). We present our ideas on a tree representation of the strategies space and finally we propose a new view of evolution of cooperation in which complexity plays a major role.

In the last sections we describe our results and we discuss about the natural behavior of this strategy and its good robustness in ecological competitions.

1.1 IPD and Artificial Life

This game is issued from the Game Theory of John von Neumann and Oskar Morgenstern and has been introduced by RAND game theorists Merrill Flood and Melvin Dresher in 1952. The idea was to introduce some *irrationality* in Game Theory, which is used as a way of modelling interactions between individuals.

This game has been found to be a very good way of studying cooperation and evolution of cooperation and thus a sort of theory of cooperation based upon reciprocity has been set in a wide literature, such as in [1, 3, 4]. The experimental studies of the IPD and its strategies need a lot of time computation and thus with the progress of computers, a lot of computer-scientists and mathematicians have studied it as they have been able to use specific methods, like genetic algorithms, on it, see [9, 2, 5, 8, 19, 21, 26, 27, 20, 31].

As cooperation is a topic of continuing interest for the social, zoological and biological sciences, a lot of works in those different fields have been made on the IPD: [6, 7, 16, 17, 18, 22, 23, 25, 29, 24].

Although all people who have studied the IPD come from different research fields, it could be said that they are all working on the same topic, which belong to the Artificial Life field, the bottom-up study of Life.

1.2 Recall on the IPD

Let two artificial agents have the choice between cooperation and defection. They play one against the other, in a synchronous manner, so that they do not know what

the other will play. They get a score according to the situation of the move:

- They both cooperate and then get both the cooperation reward, let evaluate it to R points;
- They both defect and then get both the selfish punishment, let evaluate it to P points;
- One chooses to defect while the other chooses to cooperate, then the one who has defected gets the selfish temptation salary, let it be T points, and the one who has cooperated gets the sucker score, let it be S points.

To have a dilemma, temptation must be better than cooperation, which must be better than punishment, which must be better than to be the sucker. This can be formalised as:

$$T > R > P > S$$

Since this one-shot version of the Prisoner's Dilemma is not very interesting (the most rational choice is to defect), the game is iterated, the final score being the sum of all the moves scored. Each player does not know how many moves there will be, thus each agent's strategy can be studied, to look, for instance, how each player tries to put cooperation in the game.

To avoid the one-shot Prisoner's Dilemma solution to influence strategies, by giving too much importance to temptation regarding cooperation, it is useful to add the following restriction : $2R > T + S$

With this restriction, strategies have no advantage in alternatively cooperate and defect. The classical payoff matrix used is shown in table 1.

T	5
R	3
P	1
S	0

Table 1: Classical payoff rate in the IPD

To study the behavior of strategies, two kinds of computation can be done.

The first one is a simple round-robin tournament, in which each strategy meets all other strategies. Its final score is then the sum of all scores done in each confrontation. At the end, the strategy's strength measurement is given by its range in the tournament. This is the way the *tit-for-tat* strategy has been isolated by Axelrod in [1].

The second one is a simulated ecological evolution, in which at the beginning there is a fixed population including the same quantity of each strategy. A round-robin tournament is made and then the population of bad strategies is decreased whereas good strategies obtain new elements. The simulation is repeated until the

population has been stabilised, i.e. the population does not change anymore. This is this way that nasty strategies, those who take the initiative of the first defection, have been discovered to be not very stable, because they are invaded by kind ones.

1.3 Why studying the IPD?

As we have just said in section 1.1 a lot of works have been done on the IPD and from different points of view. The common points of almost all those works are :

- they do not extensively looked for new strategies, thus good strategies and new characteristics have perhaps been missed.
- they study variations of the IPD, because they think the original model is too simple. Those variations offer more proximity to real life. Here are examples of such studies: [9, 6, 7, 16, 21, 23, 28, 20, 31].
- they seldom call into question classical results about *tit-for-tat*, which are often considered as definitive. An exception is the paper of Boyd and Loberbaum [8] showing that no pure strategy is evolutionarily stable in the IPD.

To avoid confusion with works quoted in the second point, let us call the Iterated Prisoner's Dilemma we have described, the Classical Iterated Prisoner's Dilemma (CIPD). We think that new works and new discoveries are possible on the CIPD and thus we look for good strategies for it and try to understand how and why they work. The CIPD, as a model of cooperation study, has not been totally understood. It is still a good model for this kind of problem because it is the simplest model and not everything is known about it.

One more reason, if needed, to make those researches, is that good strategies in the CIPD, are still good in a lot of variants of the game.

2 *Gradual*, a good strategy for the CIPD

2.1 Behavior of *gradual*

During our researches we have been led to test and verify classical results [1, 3]. Thus, in order to do so, we tried to create a lot of strategies to look at their behaviors and to check if they had the qualities of *tit-for-tat*. At this time we were trying to increase the efficiency of strategies by modifying their parameters little by little, and looking at the effect of the changes in round-robin tournaments and ecological simulations, until we were satisfied by the results of a strategy. During these experiments we have created many strategies including *gradual* and discovered its strength.

This strategy acts as *tit-for-tat*, except when it is time to forgive and remember the past. It uses cooperation on

the first move and then continues to do so as long as the other player cooperates. Then after the first defection of the other player, it defects one time and cooperates two times; after the second defection of the opponent, it defects two times and cooperates two times, ... after the n^{th} defection it reacts with n consecutive defections and then calms down its opponent with two cooperations.

As we can see this strategy has the same qualities as those described by Axelrod in [1] for *tit-for-tat* except one: the simplicity. *Gradual* has a memory of the game since the beginning of it.

2.2 Performance of *gradual*

We conducted some experiments with *gradual* and other strategies. Their results show the good performance of *gradual*. Other results are reported in [11].

We made two kinds of experiments with strategies: round-robin tournaments and ecological simulation. In the following example we have used 12 classical strategies. Each game consists of 1000 moves, with the classical payoff shown in table 1. The general results do not depend on the precise value of payoff and on the number of move, see [11].

2.2.1 Description of the 12 strategies used

The 12 strategies are described below.

cooperate always cooperates

defect always defects

random cooperates with a probability of 0.5

tit-for-tat cooperates on the first move and then plays what its opponent played on the previous move

spite cooperates until the opponent defects, then defects all the time

per_kind plays periodically [cooperate, cooperate, defect]

per_nasty plays periodically [defect, defect, cooperate]

soft_majo plays the opponent's most used move and cooperates in case of equality (first move considered as equality)

mistrust has the same behavior as *tit-for-tat* but defects on the first move

prober begins by playing [cooperate, defect, defect], then if the opponent cooperates on the second and the third move continues to defect, else plays *tit-for-tat*

gradual

pavlov cooperates on the first move and then cooperates only if the two players made the same move, this strategy was studied in [27].

2.2.2 Some results

Results of round-robin tournament and of a simple ecological simulation are presented in tables 2, 3 and figure 1.

They show that *gradual* clearly outperforms all other opponents and especially *tit-for-tat*, in round-robin as well as in ecological evolution. In this latter type of computation we have noticed that *gradual* and *tit-for-tat* have the same type of evolution, with the difference of quantity in favor of *gradual*, which is far away in front of all other survivors when the population is stabilised.

We notice that *gradual* is relatively strong, compared to other strategies, when opposed to clever, or probabilistic strategies like *prober* or *random*, but the most significant point is that *gradual* has never, or not often, bad scores: in almost all cases, it has been able to install cooperation with its opponents.

strategy	final score
<i>gradual</i>	33416
<i>tit-for-tat</i>	31411
soft_majo	31210
spite	30013
prober	29177
pavlov	28910
mistrust	25921
cooperate	25484
per_kind	24796
defect	24363
per_nasty	23835
random	22965

Table 3: Ordered final score in Round-robin

Numerous other experiments (with for example random subsets of strategies) confirm the conclusion that *gradual* is more robust and obtain better score than *tit-for-tat* in almost all contexts (exceptions especially constructed are obviously possible but are seldom and not easy to obtain).

A systematic study of the dynamic evolution of populations with 3 strategies confirms the conclusions of [8] that very complex phenomenon may sometimes occur in the IPD [15].

2.3 The experiment with “Pour La Science”

We have been surprised by those results and then with the help of “Pour La Science”, the french edition of the “Scientific American”, we have organized a tournament. Each reader was invited to submit his own strategy. A description of this tournament and its results can be found in [10, 12, 13]. The tournament did not used the CIPD, but a variant in which each player had the ability to give up in a definitive way with its opponent.

However the winner was a strategy which is a variant of *gradual*, adapted to the game with renunciation

	coop	def	rand	tft	spite	p_nst	p_kn	sft_mj	mist	prob	grad	pav
coop	3000	0	1481	3000	3000	999	2001	3000	2997	6	3000	3000
def	5000	1000	3003	1004	1004	2332	3668	1004	1000	1008	1340	3000
rand	4012	499	2228	2250	505	1667	2824	1980	2240	1581	940	2239
tft	3000	999	2248	3000	3000	1998	2667	3000	2500	2999	3000	3000
spite	3000	999	3010	3000	3000	2331	3663	3000	1003	1007	3000	3000
p_nst	4334	667	2502	2003	671	1666	3335	671	1999	2006	979	3002
p_kn	3666	333	2024	2667	343	1665	2334	3666	2664	2664	767	2003
sft_mj	3000	999	2380	3000	3000	2331	2001	3000	2500	2999	3000	3000
mist	3002	1000	2244	2500	1003	1999	2669	2500	1000	3000	3001	2003
prob	4996	998	2522	2999	1002	1996	2669	2999	2995	1004	2999	1998
grad	3000	915	2815	3000	3000	2219	3472	3000	2996	2999	3000	3000
pav	3000	500	2244	3000	3000	1332	2833	3000	1998	2003	3000	3000

Table 2: Round-Robin 2 by 2 scores




Figure 1: Ecological evolution

and with a different progressive function for retaliation ($N(N+1)/2$ instead of N , with N the number of opponent's previous defections). This fact makes us think that *gradual* has something *tit-for-tat* has not and that classical interpretation of *tit-for-tat* results concerning simplicity must be revisited.

2.4 Comparison with *tit-for-tat*

With the same approach to all these results as Axelrod made on *tit-for-tat*, we can say that the most three important qualities of *gradual* are: kindness (it does not begin to defect), reactivity (it defects when the opponent has defected) and forgiveness (it comes back to cooperation after punishment). These ones are well known to be *good*, but in *gradual*, unlike in *tit-for-tat*, they are not joined by simplicity. In fact the main difference between *tit-for-tat* and *gradual* is their view of the past, the former has a short one and the latter the biggest possible, since it has to know the full history of the current game

to decide what to do on the next move. We can say that *gradual* is much more *complex* than *tit-for-tat*. Let us recall an analyse of Axelrod in [1, page 110]:

The advice takes the form of four simple suggestions for how to do well in a durable Iterated Prisoner's Dilemma:

1. *Don't be envious,*
2. *Don't be the first to defect,*
3. *Reciprocate both cooperation and defection,*
4. *Don't be too clever.*

When Axelrod writes that, he states that a strategy has to be understandable, hence simple, mainly for its opponents to understand that all it wants is to establish a cooperation period. He thinks that if the strategy clearly announces how it will act, then the other player will be able to cooperate more quickly.

We have previously discussed this point in [14], for the Iterated Prisoner's Dilemma with Renunciation. We

agree with Axelrod but we think that having a clear behavior is not always a good idea, for instance against complex-clever-and-nasty or random strategies. The results of *gradual* confirm us in this opinion.

2.5 Natural inspiration of *gradual*

Finally we think that *gradual* presents a natural behavior, which we can find in our daily life, which is, for instance, used by some creditors with their debtors. Let us think of the government and the taxpayers or an electricity company and their bad customers.

This behavior can however be interpreted in two ways:

- The player is very offensive and wants to force his opponent to cooperate and thus clearly shows him that he will be more and more aggressive, so that its best choice is to cooperate.
- The player is very defensive and does not want to be exploited and thus he looks less and less for cooperation, which is risky, but plays more and more the rational choice of the single round Prisoner's Dilemma, to insure its score. However as it is not a single round game, he retries sometimes to reinstall cooperation.

Those two interpretations are two ways of looking at the game, as we have two ways of looking at our relationships with other people in real life. In the first case it tries to explain to the opponent what is the better choice for both of them, whereas in the second case it tries to protect itself. This is a kind of choice between opening or closing its relationship with the opponent, knowing none of the real made choice. It is clear that in real life this is not a simple choice to do. We think that *gradual*, instead of *tit-for-tat*, offers this type of complexity to the player and thus that Axelrod's idea about simplicity is not generally true.

2.6 Is *gradual* strong enough ?

Gradual's behavior, its performance, its inspiration of the "Pour la science" tournament winner and its natural roots, suggest that it is a strong strategy which truly outperforms *tit-for-tat*. The point that we found a strategy stronger than *tit-for-tat* and that its superiority comes from a more complex behavior is enough to think that there is maybe another more complex strategy which is better than *gradual*. *Gradual* may be improved and in order to know if it can be the case, one of the solutions is to use optimization tools on it.

3 Looking for better strategies using Genetic Algorithms

In order to try to improve *gradual's* performances, we have chosen to use a genetic algorithm to optimize *gradual*, so we have constructed a genotype and a fitness function evaluating the quality of strategies.

With this genetic algorithm we thought we would be able to run into the space of described strategies, looking for good ones, especially the ones better than *gradual*.

3.1 Genotype description

Here is the description of the genotype we have set up, using extensions of *gradual's* ideas, to describe a big family of strategies.

The extensions we have added to *gradual* are:

- ability to parameter the time it took the strategy to react at a defection;
- ability to take into account the defection made during a punishment period;
- ability to use defection, or score, threshold, to begin a punishment;
- ability to forgive randomly;
- ability to defect randomly;
- ability to modify the vision of the past.

The 19 different genes of the genotype and their meanings are given now :

begin the first move to do;

alea the number used to determine random defection;

calcul_type the type of detection used to launch a punishment period;

threshold the threshold above which punishment is launched;

forgive the number used to determine random forgiveness move;

blind the ability to update the evolution during a reaction-punishment-lull period;

vision the length of the past the strategy can see (0 being from the beginning);

punishment_evolution the type of punishment length evolution (polynomial or log.);

A_p, B_p, C_p the coefficients of punishment length evolution;

reaction_evolution the type of reaction length evolution (polynomial or log.);

A_r, B_r, C_r the coefficients of reaction length evolution;

lull_evolution the type of lull length evolution (polynomial or log.);

A_l, B_l, C_l the coefficients of lull length evolution;

Strategies described in this way play as follows: after the first move, which is **begin**, they defect with a probability of $p = 1/\mathbf{alea}$. If they do not defect then they decide if a reaction-punishment-lull period has to be launch, according to **calcul.type**, **threshold** and **vision** parameters. If such a period has to be launch, it is done with a probability of $p = 1 - 1/\mathbf{forgive}$. Parameters of such a period, *i.e.* reaction, punishment and lull length are computed according to the number of done punishments (let it be N) in one of those two manners, according, respectively, to **reaction_evolution**, **punishment_evolution**, **lull_evolution**: If the computation has to be polynomial the value will be:

$$A_i N^2 + B_i N + C_i$$

and in the logarithmic case it will be

$$A_i \frac{\log(N)}{\log(B_i)} + C_i \quad (B_i > 1)$$

with $i \in \{r, p, l\}$ and N updated at each move if it is not **blind**.

The most new parameter here is the **vision** one, which makes a big difference in the description of strategies to be used in genetic algorithms, as they have been done for now in [2, 19, 21], where strategies are limited to a small view of the past (3 or 4 moves in general).

3.2 A big space

The space of described strategies, is greatly open and depends on the way individuals are generated at the beginning of the genetic algorithms.

This space is so open that it includes almost every well known strategies:

gradual Cooperate, 0, Defection, 1, 0, Yes, 0, Polynomial, 0,1,0, Polynomial, 0,0,0, Polynomial, 0,0,2

tit-for-tat Cooperate, 0, Defection, 1, 0, Yes, 1, Polynomial, 0,0,1, Polynomial, 0,0,0, Polynomial, 0,0,0

defect Defect, 1, whatever you wish

cooperate Cooperate, 0, Defections, 0, Yes, 0, Polynomial, 0,0,0, Polynomial, 0,0,0, Polynomial, 0,0,0

With this approach, *tit-for-tat* could be seen as a degenerated strategy coming from this family and thus from *gradual*. It is a degenerated *gradual*, because some of its parameters are erased, in comparison to *gradual*.

Hence there is a chance that *gradual* could also be a degenerated strategy, coming from some X strategy, with this X better and more complex than *gradual* and then than *tit-for-tat*.

Our idea is that there is no limitation to this affiliation between good strategies and that the space of strategies is full of complex good ones, which have not been found now.

The complexity could be such that the full strategies space might be seen as a tree, where the large space of strategies described here with our genotype is only one branch of it. Maybe on another branch there are other strategies better than *tit-for-tat*, or *gradual* built with other ideas.

4 Results

4.1 How we have obtained it

The genotype we have created defines a space of 8.64×10^{15} different strategies. We have used a simple Genetic Algorithm, using uniform cross-over (*resp.* mutation) to cross over two elements of the population (*resp.* to mute one of it). The algorithm chooses randomly a vector of 19 bits, one bit for each gene and then swaps the genes of the two individuals (*resp.* mutes the gene of the individual) when there is a 1 at the gene's position in the vector.

To compute the fitness of a strategy we have used a simple idea: to be good a strategy needs to beat, in a simple round-robin tournament as those used by Axelrod, the best known strategies at this time, as well as some special ones like the simplest ones (Cooperate, Defect, Random, ...). It is however clear that this is not sufficient and that there may be more complex ways to appreciate the quality of strategies, using, for instance, its behavior in ecological simulation.

So we have chosen a selection of 34 strategies found in previous works, trying to have an heterogeneous selection of behaviors and we compute the fitness of a strategy as its rank in a round-robin tournament including the quoted 34 strategies and itself.

The initial population was made by 150 randomly generated strategies.

We have obtained some results which are easy to understand as every gene has a clearly defined contribution in the behavior of the strategy.

4.2 Remarkable values of genes

Some genes have converged to remarkable values, which give the ability to precise some of our ideas and to infirm, or confirm, some of classical ideas about qualities a strategy needs to have to be good.

The less surprising result is that in almost all our experiments, the **begin** gene has converged very quickly to the Cooperate value, which means that good strategies, we were looking for, would not be too aggressive and have great chances to be kind ones.

This idea has been encouraged by the convergence of the **alea** gene.

Like the **begin** gene, the **alea** gene has converged very quickly to the 0 value, which in our experiments is interpreted as: never defect randomly.

Since the **begin** gene has converged almost at the same time to the Cooperate value, it is clear that strate-

gies we were looking for have to be kind ones. That did not sound surprising since the idea that to be good, strategies have to be kind, is widely accepted.

Nevertheless this is a new way of confirming one of Axelrod's idea, using a different method than the one he has used to find it.

The **forgive** gene had also converged in almost all cases to the 0 value, which means that strategies must not forgive any of its opponent defection.

This is closely in relation with the fitness we have used, *i.e.* without any ecological simulation and thus push some known strategies back, like for instance the *generous tit-for-tat* family, which is however included in the strategies space described by our genotype.

The fitness we use is not dedicated to the creation of strategies which have to be strong in ecological simulation. That could be an explanation of the convergence of this gene, which eliminates some good strategies in ecological simulation.

The last remarkable convergence is the **blind** gene convergence, which has converged to the No value, which means that strategies we were running to would not forget to count opponents defection during reaction-punishment-lull periods, which is not, for instance, the case of *gradual*.

This convergence clearly shows that good strategies would be a little bit more clever, thus more complex, than *gradual* and thus more than *tit-for-tat*. This fact confirms our idea and thus goes against the generally accepted ideas about the relation between efficiency and simplicity.

The convergence of this gene is another confirmation that simplicity does not mean strength in the CIPD.

4.3 One generated strategy which outperforms *gradual*

Just for the illustration, here is the description of one of the strategies we have found with the help of the genetic algorithm we have used on our genotype:

Cooperate, 0, Defection, 1, 0, No, 5, Polynomial, 1,1,3, Polynomial, 0,0,0, Polynomial, 15,8,4

This strategy beats *gradual* and thus *tit-for-tat*, in round-robin tournament, as well as in an ecological simulation. In the two cases it has finished first just in front of *gradual*, *tit-for-tat* being two or three places behind, with a big gap in the score, or in the size of the stabilised population.

All those results have been found with the 34 strategies chosen for the fitness computation and where confirmed with an experiment involving 300 strategies.

5 Conclusions and further works

As we said, the IPD is a good model in the study of cooperation, but it has not been often studied in its classical

version. While a lot of works try to discuss, or to find how to improve the model itself, we choose to go deeper in the CIPD by looking for good strategies and by trying to understand what make those strategies strong, knowing that every results on the CIPD, can often be ported in other variations of the IPD.

We have found a strategy which leads us to reformulate some of the classical results about simplicity and then have two purposes:

1. trying to find more *complex-and-good* strategies;
2. trying to find what makes those strategies good.

For the former, we have used a genetic algorithm to look for strategies in a big space of strategies, which we have described with a genotype we have created. For the latter we just have to study the results of the experiments done by the genetic algorithm.

The first result we have found enables us to think that simplicity is not a good quality for a strategy, but that like in real-life, complexity may be advantageous.

We also think that this complexity can be so important that maybe there is not a classical linear evolution in the strength of strategies, but that the space of strategies can be seen as a tree with lots of branches. Then each branch is a large space of strategies, with good and bad ones and with some connections between branches, so that there is not a finite set of quality for a strategy to be good, but rather that particular combination of some qualities make the strength of a strategy.

We are now working on those two purposes by improving the fitness of the genetic algorithm, hoping that we could find stronger strategies; and by trying to create other genotypes, to explore other branches of the total space of strategies, hoping we could then find other combinations of qualities.

The former work implies that methods to measure the strength of a strategy have to be set, whereas the latter could reinforce our ideas about the natural emergence of the complexity in the CIPD.

A simulation software with many strategies is already available for Unix, Dos or Windows by web at <http://www.lifl.fr/~mathieu/ipd> or by anonymous ftp on the following site [ftp.lifl.fr](ftp://ftp.lifl.fr/pub/users/mathieu/soft) in `pub/users/mathieu/soft`

References

- [1] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
- [2] R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In L. Davis, editor, *Genetic Algorithms and the Simulated Annealing*, chapter 3, pages 32–41. Pitman, London, 1987.
- [3] R. Axelrod and D. Dion. The further evolution of cooperation. *Science*, 242:1385–1390, 1988.

- [4] R. Axelrod and W. D. Hamilton. The evolution of cooperation. *Science*, 211:1390–1396, 1981.
- [5] S. Banks. Exploring the foundations of artificial societies. In R. A. Brooks and P. Maes, editors, *Artificial Life, Proc. 4th International Workshop on the Synthesis and Simulation of Living Systems*, volume 4, pages 337–342. MIT Press, 1994.
- [6] J. Batali and P. Kitcher. Evolutionary dynamics of altruistic behavior in optional and compulsory versions of the iterated prisoner’s dilemma. In R. A. Brooks and P. Maes, editors, *Artificial Life, Proc. 4th International Workshop on the Synthesis and Simulation of Living Systems*, volume 4, pages 344–348. MIT Press, 1994.
- [7] J. Bendor. In good times and bad: Reciprocity in an uncertain world. *American J. of Political Science*, 31:531–558, 1987.
- [8] Robert Boyd and Jeffrey P. Lorberbaum. No pure strategy is evolutionarily stable in the repeated prisoner’s dilemma game. *Nature*, 327:58–59, 1987.
- [9] E. Ann Stanley D. Ashlock, M. D. Smucker and L. Tesfatsion. Preferential partner selection in an evolutionary study of prisoner’s dilemma. Economics R. No 35, Submitted for publication, 1994.
- [10] J.P. Delahaye. L’altruisme récompensé ? *Pour La Science*, 181:150–156, 1992.
- [11] J.P. Delahaye and P. Mathieu. Expériences sur le dilemme itéré des prisonniers. Rapport de Recherche 233, LIFL Lille CNRS (URA 369), 1992.
- [12] J.P. Delahaye and P. Mathieu. L’altruisme perfectionné. *Pour La Science*, 187:102–107, 1993.
- [13] J.P. Delahaye and P. Mathieu. L’altruisme perfectionné. Rapport de Recherche 249, LIFL Lille CNRS (URA 369), 1993.
- [14] J.P. Delahaye and P. Mathieu. Complex strategies in the iterated prisoner’s dilemma. In A. Albert, editor, *Chaos and Society*, Amsterdam, 1995. IOS Press.
- [15] J.P. Delahaye and P. Mathieu. Studies on ipd dynamics with few number of strategies : Is there any chaos in the pure iterated prisoner’s dilemma. Rr., LIFL Lille CNRS (URA 369), 1996.
- [16] M. R. Frean. The prisoner’s dilemma without synchrony. *Proc. Royal Society London*, 257(B):75–79, 1994.
- [17] H. C. J. Godfray. The evolution of forgiveness. *Nature*, 355:206–207, 1992.
- [18] N. V. Joshi. Evolution of cooperation by reciprocity within structured demes. *J. of Genetics*, 66(1):69–84, 1987.
- [19] K. Lindgren. Evolutionary phenomena in simple dynamics. In J.D. Farmer C. Langton, C. Taylor and S. Rasmissen, editors, *Proc. Workshop on Art. Life*, volume 2, pages 295–312. Santa Fe Institute, Addison Wesley, 1992.
- [20] D. Ashlock M. D. Smucker, E. Ann Stanley. Analyzing social network structures in the iterated prisoner’s dilemma with choice and refusal. RR. CS-TR-94-1259, University of Wisconsin-Madison, Département of Computer-Sciences, 1994.
- [21] C. Martino. Emergent nastiness in iterated prisoner’s dilemma games. 2.725: Design and Automation, 1995.
- [22] R. M. May. More evolution of cooperation. *Nature*, 327:15–17, 1987.
- [23] P. Molander. The optimal level of generosity in a selfish, uncertain environment. *J. of Conflict Resolution*, 29(4):611–618, 1985.
- [24] M. Nowak. Stochastic strategies in the prisoner’s dilemma. *Theoretical Population Biology*, 38:93–112, 1990.
- [25] M. Nowak and K. Sigmund. The evolution of stochastic strategies in the prisoner’s dilemma. *Acta Applicandæ Mathematicæ*, 20:247–265, 1990.
- [26] M. Nowak and K. Sigmund. Tit for tat in heterogeneous populations. *Nature*, 355:250–253, 1992.
- [27] M. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner’s dilemma game. *Nature*, 364:56–58, 1993.
- [28] M. Oliphant. Evolving cooperation in the non-iterated prisoner’s dilemma. In R. A. Brooks and P. Maes, editors, *Artificial Life, Proc. 4th International Workshop on the Synthesis and Simulation of Living Systems*, volume 4, pages 350–352. MIT Press, 1994.
- [29] R. Pool. Putting game theory to the test. *Science*, 267:1591–1593, 1995.
- [30] W. Poundstone. *Prisoner’s Dilemma : John von Neumann, Game Theory, and the Puzzle of the Bomb*. Number 0-19-286162-X. Oxford University Press, Oxford, 1993.
- [31] Xin Yao and P. J. Darwen. An experimental study of n-person iterated prisoner’s dilemma game. *Informatica*, 18:435–450, 1994.