

2009

Iterated Prisoner's Dilemma for Species

Philip Hingston
Edith Cowan University

Iterated Prisoner's Dilemma for Species

Philip Hingston, *Senior Member, IEEE*

Abstract—The Iterated Prisoner's Dilemma (IPD) is widely used to study the evolution of cooperation between self-interested agents. Existing work asks how genes that code for cooperation arise and spread through a single-species population of IPD playing agents. In this paper, we focus on competition between *different species* of agents. Making this distinction allows us to separate and examine macroevolutionary phenomena. We illustrate with some species-level simulation experiments with agents that use well-known strategies, and with species of agents that use team strategies.

I. INTRODUCTION

A beguiling puzzle of biology is how cooperative behavior can evolve in a population of selfish organisms. Ever since Axelrod and Hamilton's pioneering paper [1] in 1981, the simulated evolution of agents playing the iterated prisoner's dilemma (IPD) has been the gold standard for examining this question. IPD is a model that encapsulates the choices an organism faces regarding whether or not to cooperate with another organism, and the payoffs resulting from the choices the two organisms jointly make. A thorough coverage of the history and the main themes of this large body of research can be found in [3].

Many hundreds of papers on the topic have examined it from many directions – the effect of miscommunication, spatial models, multiple levels of cooperation, multiple players, choice of partners, signaling, selection schemes and so on. But in all these variations, the evolutionary process has been studied at the level of changing proportions of different alleles in a population of players. This level of abstraction is sometimes called “microevolution”.

In this paper, we propose a framework for studying the evolution of cooperation at the level of competition between *species*, sometimes called “macroevolution”. This is the level better suited to consider phenomena such as speciation, mutualism (cooperation between species), parallel evolution, extinction and so on.

The term “macroevolution” can be a controversial one. It is sometimes used by creationists and Intelligent Design proponents to split evolutionary theory into a part that explains variations within species (“microevolution”) and a part that explains larger scale phenomena, such as speciation (“macroevolution”). This second part then becomes a target to attack. We are not interested in these arguments here, and

we simply use “macroevolution” to refer to those phenomena of evolution that are best thought about at the level of abstraction appropriate to competition between species.

In the rest of this paper, we first introduce the iterated prisoner's dilemma and discuss related work. We then present our simulation framework, and use it to design and carry out some experiments. We begin with experiments using simple, well-known IPD strategies, to show how the framework can be used, and then move on to some more complex, successful strategies from recent IPD contests. We then consider some group strategies. We conclude with a discussion of the results of these experiments and suggest possibilities for future work.

II. ITERATED PRISONER'S DILEMMA

The Prisoner's Dilemma (PD) is a model used to study human and natural systems in which cooperation between self-interested individuals is observed or desired. It was introduced by Flood and Dresher in the early 1950's in studies applying game theory to global nuclear strategies [5]. It has also been applied to problems in psychology, economics, politics, and biology.

As PD is widely known, we refer the reader to [3] for a detailed description of the game, noting that we chose the common values $T = 5$, $R = 3$, $P = 1$, and $S = 0$.

A case by case analysis shows that the best way for a self-interested player to play PD is always to defect, no matter what the other player does. This leads us to consider the Iterated Prisoner's Dilemma (IPD), in which the players play a sequence of games of PD against each other.

In IPD, player strategies are rules that determine (perhaps stochastically) a player's next move in any given game situation (which can include the history of the game to that point). Each player's aim is to maximize his total payoff over the series. To prevent players anticipating the end of the series (which again leads to mutual defection), the series continues with some fixed probability, called the *discount factor* (so called because in a mathematical sense this is equivalent to an infinitely repeated game where future payoffs are discounted).

Some well-known IPD strategies are:

- TitForTat: cooperate on the first move, and play the opponent's previous move after that;
- Grim: cooperate on the first move, and keep cooperating unless the opponent defects, in which case, defect forever;

Philip Hingston is an Associate Professor of Computer Science with the School of Computer and Security Science, Edith Cowan University (e-mail: phi@philiphingston.com).

- Pavlov: cooperate on the first move, and on subsequent moves, switch strategies if you were punished on the previous move.

Many variations of IPD have been studied, using many approaches. Variations include different classes of strategies, noisy moves, noisy payoffs, alternating non-simultaneous moves, signalling and so on. Approaches used include game theory, evolutionary methods, and machine learning. See [3] for a nice review covering work up to 2007.

A. IPD Tournaments

Around 1980, Robert Axelrod staged two round-robin “tournaments” between computer programs designed by participants to play IPD. Many sophisticated programs were submitted. In each case, the winner was Anil Rapaport’s submission, a program that simply played TitForTat. In 1987, Axelrod carried out computer simulations using a genetic algorithm to evolve populations of strategies playing the IPD against each other [1]. In these simulations, TitForTat-like strategies often arose, but other, more complicated strategies sometimes evolved that outperformed TitForTat in particular populations. Axelrod used this to illustrate that there is no “best” strategy for playing the IPD in such an evolving population, because success depends on the mix of other strategies present in the population.

Over the years, other tournaments have been held, including a series of competitions held in 2004 at the IEEE Congress on Evolutionary Computation, to celebrate the 20th anniversary of Axelrod’s 1984 book on the subject [2]. Controversially, competitors were allowed to enter more than one strategy, and some took advantage of this to enter strategies that colluded with each other – so called group strategies – in order to improve the chances of one of their number of winning the competition. The competitions were repeated at the 2005 IEEE Symposium on Computational Intelligence and Games (CIG’05), along with an extra competition in which group strategies were not allowed.

III. RELATED WORK

There are many existing examples, often ad hoc, of evolutionary simulations at the *species level*. A famous one is Lovelock’s Daisyworld [12], which models the interaction between two species of daisy and their effect on climate, illustrating the Gaia hypothesis. A popular subject for simulation is the predator-prey relationship: see, for example [11]. In an example closely related to our topic, Rankin et al [9] used a multispecies simulation to show that in a competition between species, the more selfish species tend to go extinct – at least in the scenarios they simulated (but note that their study did not use IPD as the vehicle for representing cooperative versus selfish behavior).

Despite the common occurrence of species-level evolutionary simulations, we are not aware of any other species-level simulation studies focused on IPD. Many existing studies do sometimes interpret specific alleles as “species”, and refer to their “extinction”, for example.

However, alleles don’t really behave like species except in an allegorical sense. The essential property of species that is missing is that species are reproductively isolated. For example, “extinction” is but a temporary setback for an allele – mutation can always bring it back to life. Similarly, a single mutation changes one allele to another, whereas true speciation is much more complex. An interesting contribution in this area is [4], where fitness sharing is used to encourage specialization of a population into species. There have been other IPD studies where restricted reproductive choices are used in a similar way, simulating the reproductive isolation of species to some degree. There is even the suggestion that IPD-like competition can help speciation in Nature, see, for example, [7].

Our experiments on group strategies were obviously motivated by the group strategies created for the 2004/5 IPD Competitions. Another approach to group-aware strategies is the work involving the use of tag systems to allow agents to recognize different player types, for example [10]. The approach we use here is more straightforward, but does not support study of issues such as mimicry and signaling.

IV. A SIMULATION FRAMEWORK

Evolution in Nature can be thought of as a competition between organisms using different survival strategies. In this work, we want to examine one specific aspect of evolution, so we need to lay out the rules of the competition.

Informally, here are our rules:

- We are interested in a population of IPD playing agents;
- The population consists of agents belonging to a number of distinct species;
- Each species has its own specific kind of genome;
- Agents reproduce asexually, with the child’s genes derived from its parents by mutation - the child is the same species as its parent;
- An agent’s genes determine the strategy the agent uses when playing IPD;
- An agent’s fitness is determined as the average payoff it receives from playing IPD games against the rest of the population;
- An agent’s reproductive success is proportional to its fitness.

The outcome of this competition depends on many factors: chance, population size, the species and their initial proportions in the population, as well as on how, specifically, the rules of the competition are operationalised. By manipulating some of these, we aim to understand better how cooperative behaviour can evolve in populations of self-interested individuals, and what factors affect that evolution.

We implemented this framework in Java using a population is made up of Organisms, each containing a Genotype and a Phenotype. The Genotype determines the species of the Organism. The Phenotype (determined from the Genotype) is the agent’s strategy.

Although the framework provides for mutation, in the experiments reported here, we keep it simple by making mutation a no-op. Thus we are focused, in this initial exploration, on competition between species, and ignore competition within species.

```
public class Organism
{
    public Genotype genotype;
    public Phenotype phenotype;
    public double fitness;
}

public interface Genotype
{
    public Genotype copy();
    public void mutate();
    public Phenotype develop();
}

public interface Phenotype
{
    public Move getFirstMove();
    public Move getNextMove(
        Move oppLastMove);
}
```

We simulate the evolution of a population of Organisms as shown in the pseudo-code below:

Inputs: Initial number of organisms of each species

1. Create initial population of Organisms
2. While not done do
3. For each pair of Organisms O1 and O2
4. Play O1 against O2 in a game of IPD
5. O1.fitness += O1's total payoff
6. O2.fitness += O2's total payoff
7. Start a new population
8. While new population not complete
9. Select a parent O
10. C = O.genotype.copy()
11. C.mutate()
12. P = C.develop()
13. Add a new Organism(C, P, 0) to the population
14. End While
15. End While

Notice that in the case where there is only one species, this framework reduces to the more usual kind of evolutionary simulation.

A. Example – AllC, AllD and TitForTat

We use this framework to run experiments with different kinds of IPD strategies, to examine the species-level phenomena that we can observe. To illustrate, we first examine some simple, well-known strategies, AllC – where the agent ignores its opponent and cooperates at all times, AllD – where the agent always defects, and TitForTat. We start with 20 individuals in each species. The discount rate is 0.98, giving an average game length of 50 moves. We use

roulette-wheel selection. Each simulation is run for 100 generations.

During the course of a simulation, it is possible for a species to go extinct. This is a key point of departure from typical evolutionary simulations, where all the agents are of the same species, and interest is in the changing proportions of different individual-level strategies (we refer to these as individual-level strategies even though they may involve interactions between different population sub-groups). These strategies never really go extinct – they can always make a comeback through mutation or crossover. In contrast, we are examining species-level interactions, where extinction is forever.

In a sense, the real competition between species is to see who does best at survival in the long term – a short-term drop in population numbers doesn't necessarily spell failure. However, if the population of a species drops too low, then an unlucky sequence of events might lead to extinction.

There are at least three different kinds of outcomes for these simulations. One possibility is that only one species survives for the whole 100 generations. In this case, we call that species the *sole survivor*. Another possibility is that more than one species lasts 100 generations. It might be that the surviving species have reached some kind of equilibrium or limit cycle, or it might be that one would be sole survivor if we ran the simulation longer. It is difficult for us to tell the difference between these last two cases.

As the outcome of any single simulation is affected by chance, we run 100 simulations, and gather some statistics on the outcomes in terms of extinctions and sole survivors. Table 1 below shows the results for the present example.

TABLE 1
EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 20 ALLC, 20 ALLD, AND 20 TITFORTAT

	AllC	AllD	TitForTat
extinction	83	99	11
sole survivor	10	1	82

Here we see that AllD almost always goes extinct (99%), while AllC usually does (83%), and TitForTat only occasionally does (11%). It is very rare (1%) for AllD to be sole survivor. TitForTat is usually the sole survivor (82%). We can see what happens in a little more detail with a plot showing the mean population sizes at each generation over the 100 runs of the simulation, as in Figure 1.

The plot does not show a typical run – the course of each individual run may be quite different – but we can see some general features. The number of AllD initially rises, on average, while AllC falls. TitForTat rises also, eventually suppressing the AllD population. By the time AllD is eliminated, TitForTat exists in higher numbers than AllC. At this point, the population is playing almost exclusively cooperatively (100% in those runs where AllD is extinct), so the competition becomes a kind of “random walk”, where everyone's fitness is more or less equal, and stochastic variation causes the relative proportions to wander a little.

This kind of result has been seen in many previous studies. In our species-based simulations, an unlucky sequence of wanderings will result in either AIID or TitForTat going extinct if we wait long enough, which happens in all but 7 simulations in this case.




Figure 1 - Mean numbers in each generation, starting with 20 AIIC, 20 AIID, 20 TitForTat

If we increase the initial population to 50 of each species, we get the results shown in Table 2 below.

TABLE 2
EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 50 AIIC, 50 AIID, AND 50 TITFORTAT

	AIIC	AIID	TitForTat
extinction	54	100	0
sole survivor	0	0	54

We see that the stochastic effects are less with a larger total population. TitForTat always survives, and AIID always goes extinct. AIIC goes extinct 54% of the time, while AIIC and TitForTat both survive the other 46% of the time. While we have omitted the mean numbers plot here, it is similar to Figure 1, with TitForTat having an average final population of about 125 to AIIC's 25, a ratio of about 5:1, compared with about 3:1 for the smaller population case.

TABLE 3
EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 100 AIIC, 100 AIID, AND 100 TITFORTAT

	AIIC	AIID	TitForTat
extinction	31	100	0
sole survivor	0	0	31

With starting populations of 100, the results are qualitatively similar, except that AIIC only goes extinct 31% of the time after 100 generations. The final ratio of TitForTat to AIIC is about 5:1, the same ratio as for 50 initial agents per species.

V. COMPETITION ENTRIES

In the CIG 2005 IPD Competitions, Competition 4 was a re-run of Axelrod's original competition. In particular, only one entry was allowed per contestant, and group strategies were not allowed. The form of the competition was similar to a single generation of one of our simulations – fitness levels at the end of the first generation corresponded to the final scores in the competition. What would happen if we were to

continue the competition into successive generations? Arguable, this would test how well the strategies used in the competition would fare in an evolutionary context. Of course, these strategies were not designed to be used in this way, but we were interested to see the results.

Therefore, we present here an experiment in which we use some of the competition entries and subject them to our simulation framework. For practical reasons, we have chosen to include only the top 4 entries, along with the "standard" strategy Rand, a random player, that was included in the competition by the organisers. The entered strategies, in the order they finished, are: Adaptive Pavlov (our own implementation, based on our best understanding of [8]), Omega TitForTat (the actual implementation from the competition), Modeller (modified for a 0.98 discount rate), and Gradual (the competition implementation). We started with 20 agents of each species. Extinction and survival results are shown in Table 4.

TABLE 4
EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 10 AGENTS FROM EACH OF 13 COMPETITION STRATEGIES.

	extinct	sole survivor
APavlov	61	4
OTFT	46	10
Modeller	47	5
Gradual	52	6
Rand	100	0

Thus Omega TitForTat is the winner of the species-level competition. The mean numbers plot, Figure 2, shows that Rand is quickly eliminated, and the other 4 strategies are closely matched.

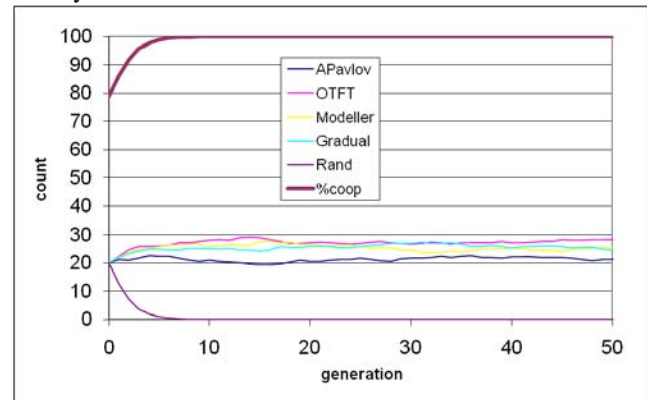


Figure 2 - Mean numbers in each generation, starting with 20 APavlov, 20 OTFT, 20 Modeller, 20 Gradual and 20 Rand

The reason APavlov does not win here may be the different mix of opponents (there were 50 entrants in the original competition), or it may be that our implementation of the strategy is not identical to the competition one (the description is ambiguous on a couple of points).

Individual runs of the simulation look quite different from Figure 2. We give an example in Figure 3. On this run, Modeller was doing best in the first couple of generations, but Adaptive Pavlov caught a lucky break in generation 3. These two continued to be favoured by chance, with the

other strategies dying out by generation 25.




Figure 3 - Numbers in each generation on a single run, starting with 20 APavlov, 20 OTFT, 20 Modeller, 20 Gradual and 20 Rand

VI. SOME GROUP STRATEGIES

So far we have looked at agents that do not distinguish between opponents – that is, agents that consider all opponents to be identical at the start of each new IPD game. In this section, we ask: what changes if agents are able to identify the species of the opponent. Such an ability seems biologically reasonable for studying species-level evolution, and would also be appropriate for non-biological scenarios in which we are interested in group strategies.

Group strategies proved very successful under the rules of the 2004/5 IPD Competitions, in those competitions where group entries were allowed. Agents used special sequences in the first few moves of the game to allow members of the same group to recognise each other, and then to play differently against different players. While we could do the same here, it seems simpler to provide all agents with the innate ability to identify an opponent’s species.

In Nature, it would often be reasonable to assume that organisms can identify the species of another organism, as long as we put aside signaling and mimicry, interesting as they are.

We can accommodate group strategies by modifying our framework as follows, adding the opponent’s genotype class as another argument to the Phenotype methods:

```
public interface Phenotype
{
    ...
    public Move getFirstMove(Class oClass);
    public Move getNextMove(
        Class oClass,
        Move oppLastMove);
}
```

Many different strategies have been developed for IPD, and the possibilities for group strategies are no doubt just as numerous. In the rest of this section, we consider just a couple of possible strategies and begin to explore the complexities of their interactions.

A. Master/Slave Group Strategies

The group strategies entered in the 2004/5 IPD Competitions work roughly as follows. There are two types of agents in a team: Master and Slave. Slaves sacrifice themselves for Masters, by repeatedly cooperating, allowing the Master to constantly defect, giving the Master the maximum payoff. A Slave playing another Slave cooperates. A Slave playing any player not in its team always defects, preventing the other player from getting a good payoff. A Master playing another Master cooperates, maximizing their joint payoff. A Master playing any player not in its team plays TitForTat (or some other good strategy), so as to maximize its own payoff against other players. So, for example, the Master’s Phenotype uses this method:

```
public Move getNextMove(Class oClass,
                        Move oppLastMove)
{
    if(oClass == Master.class)
        return COOPERATE;
    else if(oClass == Slave.class)
        return DEFECT;
    else return oppLastMove;
}
```

These strategies were designed for the context of the competitions: a good strategy was to use one Master and as many Slaves as the competition allowed. The aim was for the Master to do well, and the poor outcome for the Slaves was of no consequence. In the context of an evolutionary contest, the likely consequence for the Slaves is extinction of their species, and it is not clear whether there would be any lasting benefit for the Masters.

An alternative way to model the Master/Slave strategy would be to make Master and Slave different roles within a single species. The assignment of roles could be decided either socially (for example, individuals could switch between roles if there seem to be too many of one and not enough of the other), or genetically. Both of these would be interesting, but we have not yet attempted either.

B. The Clique Strategy

A commonly observed group strategy among humans is that of the clique, in which members cooperate only with other members of the clique. Here we model this behaviour by creating a species that uses this strategy. So, for example, one method of the Clique phenotype would be:

```
public Move getNextMove(Class oClass,
                        Move oppLastMove)
{
    if(oClass == Clique.class)
        return COOPERATE;
    else return DEFECT;
}
```

This strategy tries to give maximum assistance to clique members while denying succor to outsiders. An apparent weakness is its inability to get a good reward from outsiders

who cannot be exploited by defection, like TitForTat.

C. Simulation results

We begin with some experiments to explore the Master/Slave strategy, by pitting some familiar strategies against a Master/Slave team. Let's start with AllC.

1) Experiment C.1 – Master/Slave versus AllC

In this experiment, we started with 50 AllC agents, 25 Masters and 25 Slaves. As the Master/Slave agents are acting as a team, we allow 50 agents to be split between Masters and Slaves. Table 5 gives the extinction/survival results:

TABLE 5
EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 25 MASTER, 25 SLAVE, AND 50 ALLC

	Master	Slave	AllC
extinction	3	100	71
sole.survivor	71	0	3

As we see, the Slaves always go extinct, Masters usually (71%) become the sole surviving species, sometimes (26%) Masters and AllC both survive, and rarely (3%) AllC is the sole survivor. A more detailed picture is given by the Mean numbers plot, Figure 4 below.

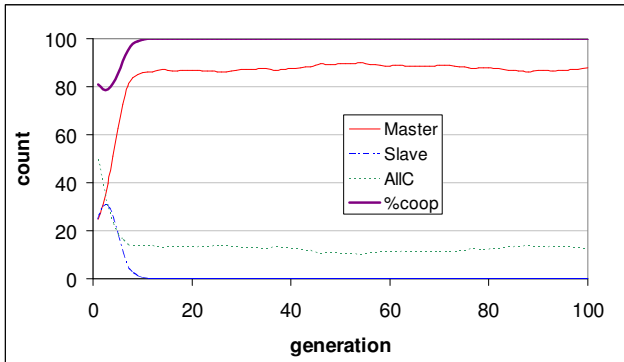


Figure 4 – Mean numbers in each generation, starting with 25 Master, 25 Slave, 50 AllC

The Slaves sacrifice themselves in the first few generations (generally going extinct at about generation 10) to give the Masters and initial advantage over AllC. By that time Masters make up about 88% of the population, and both surviving species will cooperate 100% of the time from then on. Stochastic variation then determines which species, if any, goes extinct by generation 100.

Our choice of 25 Masters and 25 Slaves was arbitrary. Would the team do better with more Slaves and fewer Masters? Perhaps starting with too few Masters would leave them still in the minority at the point when the Slaves go extinct? The next experiment examines this.

2) Experiment C.2 – Master/Slave 5/45 versus AllC

In this experiment, we start with 5 Masters, 45 Slaves, and 50 AllC. Table 6 gives the extinction/survival results. We see that the team as a whole is more successful than before, with AllC going extinct in every run. Surprisingly, the Slaves actually out-competed the Masters in one run. How did this happen? The Mean numbers plot, Figure 5, offers an explanation.

TABLE 6
EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 5 MASTER, 45 SLAVE, AND 50 ALLC

	Master	Slave	AllC
extinction	1	99	100
sole.survivor	99	1	0




Figure 5 - Mean numbers in each generation, starting with 5 Master, 45 Slave, 50 AllC

This time, the Slaves are seen to rapidly rise to nearly 80% of the population, by which time AllC is almost eliminated. By the time AllC goes extinct, the Masters have increased to more than 20% of the population, and are usually (99%) able to assert themselves and take over from the Slaves.

What happens if we start with more Masters? If we start with 45 Masters and 5 Slaves, then the Masters are sole survivors about 34% of the time, AllC about 17%, and the rest of the time, Masters and AllC both survive (with about 66% Masters and 34% AllC). As expected, a smaller initial number of Masters is more effective.

3) Experiment C.3- Master/Slave versus TitForTat

Having disposed of AllC, we now ask how the Master/Slave strategy might do against a more difficult opponent, say TitForTat. Table 7 below gives the extinction/survival results for 100 runs starting with 5 Masters, 45 Slaves, and 50 TitForTat agents.

TABLE 7
EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 5 MASTER, 45 SLAVE, AND 50 TITFOR TAT

	Master	Slave	TitForTat
extinction	14	99	45
sole.survivor	44	1	13

Here we see that the Master/Slave team is also effective against TitForTat, driving TitForTat to extinction almost half the time (45%), and going extinct only 13% of the time. The mean numbers plot, Figure 6, gives more detail.

The Slaves drive down TitForTat to around 35% on average by the time they go extinct in about generation 10. As with the competitions against AllC, 100% cooperation then ensues, with drift causing extinction of either Master or TitForTat by generation 100 about 58% of the time.




Figure 6 - Mean numbers in each generation, starting with 5 Master, 45 Slave, 50 TitForTat

4) Experiment C.4- Clique versus AllC

Starting with 50 Clique and 50 AllC agents, Clique drives AllC to extinction in about 5 generations. This is not surprising given AllC's lack of defense against exploitation.

5) Experiment C.5 – Clique versus TitForTat

In this experiment we pit 50 Clique against 50 TitForTat agents. Extinction and survival rates are given in Table 8.

TABLE 8

EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 50 CLIQUE AND 50 TITFOR TAT

	Clique	TitForTat
extinction	42	58
sole survivor	58	42

Clique has a small edge against TitForTat when they start with equal numbers. As Figure 7 shows, one species or the other is extinct by about generation 20 (we know this because we see 100% cooperation).




Figure 7 - Mean numbers in each generation, starting with 50 Clique and 50 TitForTat

Clique gains its advantage over TitForTat by defecting on the first move of the game. The average payoffs in games between Clique and TitForTat (with a discount of 0.98) is 54/50 for Clique versus 49/50 for TitForTat. This gives equal average payoffs over one generation when the proportion of Clique agents in the population is about 48.73%. If we start the simulation with 49 Clique and 51 TitForTat, Clique is the sole survivor about 51% of the time. Starting with 48 Clique and 52 TitForTat, then Clique is the sole survivor only about 35% of the time.

6) Experiment C.6 – Master/Slave versus Clique

In this experiment the two group strategies face off. Perhaps surprisingly, Clique is sole survivor 62% of the time, compared with Master 38% of the time.

TABLE 9

EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 5 MASTER, 45 SLAVE, AND 50 CLIQUE

	Master	Slave	Clique
extinction	62	100	38
sole survivor	38	0	62




Figure 8 - Mean numbers in each generation, starting with 5 Master, 45 Slave, and 50 Clique

The mean numbers plot suggests an explanation. Comparing Figure 6 and Figure 8, we see that Clique is not so damaged by Slave as TitForTat is, so that Master, starting with low numbers, never manages to catch up to Clique.

7) Experiment C.7 – All in

The success of an IPD strategy depends on the population it is part of, and it's this that makes the dynamics of an evolutionary simulation of IPD so rich and complex. In this experiment we throw all our strategies together: 5 Masters, 45 Slaves, 50 Clique, 50 TitForTat and 50 AllC.

TABLE 10

EXTINCTION AND SURVIVAL RESULTS FOR 100 RUNS STARTING WITH 5 MASTER, 45 SLAVE, 50 CLIQUE, 50 TITFOR TAT AND 50 ALLC

	Master	Slave	Clique	TitForTat	AllC
extinction	48	100	53	67	97
sole survivor	19	0	47	1	0

In avoiding extinction, the Master/Slave team just shades Clique, while Clique was best at eliminating the competition – the Clique strategy either goes extinct itself or drives the other strategies extinct – no live and let live here. In Figure 9, we show the mean numbers for each species in the first 20 generations, before full cooperation is established.

We see that Clique and Master both increase initially, with Clique getting a slightly faster start. If Clique's start is fast enough, and it reaches a large enough proportion of the population, it ruthlessly eliminates the competition. If not, then the hard to exploit Master and TitForTat strategies are together able to hold it off, and Master generally prevails by virtue of having a larger proportion of the population at the point when all the surviving species start cooperating.

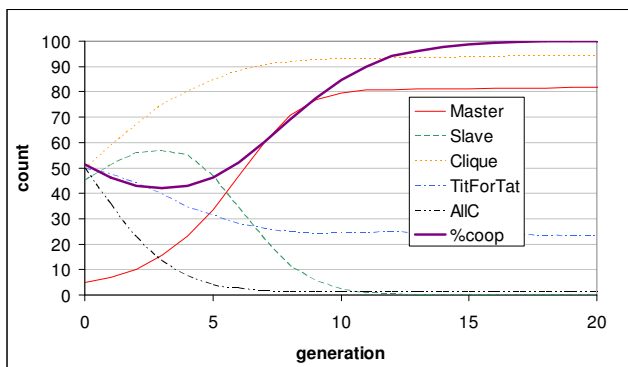


Figure 9 - Mean numbers in each generation, starting with 5 Master, 45 Slave, 50 Clique, 50 TitForTat and 50 AIC

It is perhaps surprising that the very simple Clique strategy is such a strong contender in these contests, but we must remember that the outcome is very dependent on the initial composition of the population.

VII. DISCUSSION AND FUTURE WORK

These experiments illustrate that the phenomena of species-level evolution of cooperation are complex and different from those observed in single-species simulations.

The experiments in subsection IV.A revisited the familiar case of competition between naïve cooperation, naïve defection and cooperating but non-exploitable strategies like TitForTat. The results are similar, with defection successful initially, then TitForTat punishing the defectors, allowing cooperation to re-establish. The difference in this case is that the defectors are a separate species, which goes extinct, leaving cooperators and TitForTat at the same fitness level, until one or the other drifts into extinction.

A similar story is seen in section V, where clearly inferior species inevitably go extinct, and where the relative success of well-matched species is expressed in different extinction and survival rates over many simulation runs.

The strategies in those experiments are all adaptive to some degree, with the Modeller strategy building an explicit opponent model. But they adapt only to each individual opponent during the course of a game. None of this learning carries over between opponents. Imagine instead a strategy that collects population statistics over a number of games against different opponents, and makes use of this information to more quickly identify an opponent's likely strategy. This might allow agents to learn the characteristics of the other species in the population, for example. The evolutionary process itself provides the means for another kind of adaptation, at the species level. Species that mutate can adapt their strategies in evolutionary time, by means of mutation and selection, to respond to changes in the composition of the population. These two levels of adaptation would support studies of the interaction of learning and evolution along the lines of [6].

In section VI, we considered group strategies, using species to represent groups. We found that there are group strategies that can do better than non-group strategies in

some circumstances. The Master/Slave combination that was used in some of the 2004/5 IPD Competitions defeats TitForTat in a head-on contest, for example, as does the unobtrusive Clique strategy that we introduced. A question yet to be explored is how such strategies might evolve.

Our framework could be used to model various other macroevolutionary phenomena, such as invasion by a separately evolved species, coextinction and so on. However it does not provide for speciation. This would be an interesting extension.

VIII. CONCLUSIONS

We have introduced a novel framework for studying the IPD in the context of species-level evolution.

Within this framework, it is simple to devise experiments to explore specific aspects or phenomena of macroevolution. Some experiments have been introduced to illustrate. Using survival and extinction as the yardstick, interactions between some simple, well-known strategies were examined. We then re-examined some more complex strategies from the 2005 IPD Competition. Finally, we experimented with group strategies, using species to represent groups. We believe that this framework provides a flexible and economical platform for simulating species-level evolution.

REFERENCES

- [1] Axelrod, R. and Hamilton, W. The Evolution of Cooperation, Science, Vol 211, Issue 4489, 1390-1396, 1981.
- [2] Axelrod, R. The Evolution of Cooperation, BASIC Books, New York, 1984.
- [3] Chong, S.Y., Humble, J., Kendall, G., Li, J. and Yao, X. Iterated Prisoner's Dilemma and Evolutionary Game Theory, in X. Yao, G. Kendall and S.Y.Chong Eds. The Iterated Prisoner's Dilemma: Celebrating the 20th Anniversary, World Scientific, 23-62, 2007.
- [4] Darwin, P.J. and Yao, X. Speciation as Automatic Categorical Modularisation, IEEE Transactions on Evolutionary Computation, Volume 1, Issue 2, 101-108, 1997.
- [5] Flood, M. Some experimental games, Research Memorandum RM-789, RAND Corporation, Santa Monica, CA, 1952.
- [6] Hingston, P., Dyer, D., Barone, L., French, T. and Kendall, G. Opponent Modelling, Evolution, and the Iterated Prisoner's Dilemma, in X. Yao, G. Kendall and S.Y.Chong Eds. The Iterated Prisoner's Dilemma: Celebrating the 20th Anniversary, World Scientific, 139-170, 2007.
- [7] Hochberg, M., Sinervo, B. and Brown, S. Socially Mediated Speciation. Evolution, Volume 57, Issue 1, 154-158, 2007.
- [8] Li, J. How to Design a Strategy to Win an IPD Tournament, in X. Yao, G. Kendall and S.Y.Chong Eds. The Iterated Prisoner's Dilemma: Celebrating the 20th Anniversary, World Scientific, 89-104, 2007.
- [9] Rankin, D.J., Lopez-Sepulcre, Foster, K.R. and Kokko, H. Species-level selection reduces selfishness through competitive exclusion. Journal of Evolutionary Biology, Vol. 20, Issue 4, 1459-1468, 2007.
- [10] Riolo, R. The Effects of Tag-mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoners Dilemma. Nature 414, pages 441-443, 2000.
- [11] Schmickl, T. and Crailsheim, K. Bubbleworld.Evo: Artificial Evolution of Behavioral Decisions in a Simulated Predator-Prey Ecosystem in From Animals to Animats 9, 594-605, LNCS 4095, 2006.
- [12] Watson, A.J. and Lovelock, J.E. Biological homeostasis of the global environment: the parable of Daisyworld. Tellus B (International Meteorological Institute) 35 (4): 286-9, 1983.