

Evolution of Altruism in Optional and Compulsory Games

JOHN BATALI[†] AND PHILIP KITCHER[‡]

*Departments of [†]Cognitive Science and [‡]Philosophy, University of California at San Diego,
La Jolla, CA 92093-0302, U.S.A.*

(Received on 14 September 1993, Accepted in revised form on 3 January 1995)

In “optional” variants of the iterated prisoner’s dilemma, players may choose whether or not to participate. Members of evolving populations playing optional variants of the iterated prisoner’s dilemma by following inherited strategies tend to cooperate more than do members of populations playing the standard, “compulsory” version. This result is due to dynamical properties of the evolving systems: the populations playing the compulsory game can become stuck in states of low cooperation that last many generations, while the optional games provide routes out of such states to states of high cooperation. Computational simulations of the evolution of populations playing these games support these analytic results and illustrate the interactions between the genetic representation of strategies and the composition of populations in which those strategies are deployed.

© 1995 Academic Press Limited

1. Introduction

Contemporary biological discussions of the evolution of altruism define altruistic behavior as that which increases the fitness of another animal at a cost in fitness to the animal engaging in the behavior. In the past decades, study of the evolution of genetic dispositions to altruistic behavior in this sense has been advanced by considering two special instances: (i) cases in which the animal’s own reproductive losses are made up through increases in the reproductive successes of kin, and (ii) cases in which the short-term losses of altruistic actions are made up through a system of reciprocation. Separating these instances is heuristically helpful, but artificial, since it is possible that interactions with relatives might involve short-term fitness losses that are made up through contributions from both sources.

Our concern is with cases of type (ii) and with a modification of what has become the standard way of dealing with such cases. Reciprocal altruism was originally introduced in a seminal paper by Robert Trivers (1971). Following the work of Axelrod and

Hamilton (Axelrod & Hamilton, 1981; Axelrod, 1984), reciprocal altruism has been explored by considering strategies for playing iterated Prisoner’s Dilemma (henceforth PD). Initially it appeared that the strategy *Tit-for-Tat* would be evolutionarily stable (in the sense of Maynard Smith & Price, 1973; see Maynard-Smith, 1982), but further investigation has shown that it is not so (Boyd & Lorberbaum, 1987; Farrell & Ware, 1989; Mesterton-Gibbons, 1992). Indeed, the many discussions of the Axelrod–Hamilton approach have revealed unsuspected complexities in the selection of strategies for playing the iterated Prisoner’s Dilemma (Dugatkin & Wilson, 1991; Peck & Feldman, 1986; Lindgren 1992). Our aim in this paper is to suggest a modification of the Axelrod–Hamilton scenario, and to explore the dynamics of the selection of altruistic strategies.

We begin by presenting a new version of the Prisoner’s Dilemma in which players have the option of not participating in interactions. Analytical evaluations of strategies for this game are presented to show that populations playing the optional games will achieve states of high cooperation more reliably than will populations playing the standard, compulsory, Prisoner’s Dilemma. We then present the results of

E-mail: batali@cogsci.ucsd.edu or pkitcher@ucsd.edu

computer simulations of these games that confirm the analytic results, and also illustrate the dynamics of strategy changes in different kinds of game.

1.1. COMPULSORY AND OPTIONAL PRISONER'S DILEMMAS

In each round of the standard Prisoner's Dilemma, each of two players must choose one of two actions: C, (Cooperate) or D (Defect). After each player has chosen, a payoff for each is computed from a payoff matrix as follows:

	C	D
C	R	T
D	S	P

with R = reward for mutual cooperation; T = "traitor" payoff to defector if other cooperates; S = "sucker" payoff to cooperator if other defects; P = payoff for mutual defection, where $T > R > P > S$ and $2R > T + S$.

This standard game is compulsory in the following sense: first, the players have no choice of partner but are forced to interact with the assigned individual; second, they have no choice but to play: there is no individual asocial behavior available to them. In focusing on the iterated Prisoner's Dilemma, Axelrod and Hamilton considered the evolution of altruism in a population of animals that is forced into social activity, so that the problem is posed in terms of the victory of cooperation over anti-social behavior. As we shall explain below, we believe that it may be more realistic to consider the evolution of altruism in situations in which the possibility of asociality is present, and that this will make a difference to the evolutionary scenarios.

Various authors have considered the possibility that animals facing repeated game-like situations might either have an asocial option (Miller, 1967; Lima, 1989) or be able to discriminate partners (Nöe *et al.*, 1991; Peck, 1990; Bull & Rice, 1991), but we believe that the concept of an optional game has not previously been clearly articulated (the closest approach we know is that of Stanley *et al.*, 1994). In an optional game, individuals have the possibility of signaling willingness to play to other individuals in the population. When signals of mutual willingness are given, the individuals play the game together. An animal that is unwilling to play with any of the animals willing to play with it is forced to opt out, performing some asocial behavior. We allow that such asocial behavior might have some intermediate benefit to the player, which we represent by the payoff W . In the cases of interest, $T > R > W > P > S$, i.e. the opting-out payoff is less advantageous to both players

than is mutual cooperation, but is more advantageous to both players than is mutual defection. Optional games of this sort will be referred to as "fully optional" versions of the iterated Prisoner's Dilemma.

Intermediate between fully optional and compulsory games are "semi-optional" games. These are of two types. In one version, individuals have no choice of partners, but they do have the ability to opt out. In the other, they cannot opt out, but can signal willingness to play with particular partners. The latter type of game will be of no concern to us, and we shall henceforth use "semi-optional" game to refer only to those games in which there is the possibility of opting out, but no choice of partners. We hope that the distinction between compulsory, fully optional, and semi-optional games is now sufficiently clear, and proceed with an illustrative example.

1.2. AN OPTIONAL-GAME MODEL OF GROOMING

Mutual grooming among primates serves the function of removing parasites from the fur. An animal that grooms another incurs a cost, depending on the length of time taken up in grooming, an animal that is groomed receives a benefit, which we can also assume to be proportional to the amount of time spent in grooming. Assume that there are two possibilities for each animal: to groom for a long period (C) or to groom for a short period (D). Let the cost of grooming for a short period be c_0 , the benefit from a short period of grooming be b_0 , the cost of a long period of grooming be c_2 , and the benefit from a long period of grooming be b_2 . Then, in a single interaction between two animals, the payoff matrix will be

	C	D
C	$b_2 - c_2$	$b_0 - c_2$
D	$b_2 - c_0$	$b_0 - c_0$

Since $b_2 > b_0$ and $c_2 > c_0$, it is clear that $b_2 - c_0$ is the highest payoff and $b_0 - c_2$ is the lowest. We may reasonably assume that a more thorough grooming is worth the extra time spent, and thus that $b_2 - c_2 > b_0 - c_0$. Given this inequality, it follows that $2(b_2 - c_2) > b_2 - c_2 + b_0 - c_0 = (b_2 - c_0) + (b_0 - c_2)$. Thus the payoff matrix may be rewritten as:

	C	D
C	R	T
D	S	P

with: $R = b_2 - c_2$; $T = b_2 - c_0$; $S = b_0 - c_2$; $P = b_0 - c_0$; and so $T > R > P > S$ and $2R > (T + S)$. These are the conditions for a compulsory Prisoner's Dilemma.

But our primates surely have another option—they can groom themselves. Presumably the payoff from self-grooming is less than one would receive from an earnest grooming job from another, but greater than the payoff from a more desultory performance. This can be modeled by assigning an intermediate cost c_1 , and an intermediate benefit b_1 for self-grooming. The payoff from self-grooming is thus $b_1 - c_1$, and we will assume that $b_2 - c_2 > b_1 - c_1 > b_0 - c_0$. Given that primates can also signal willingness to engage in grooming interactions with some animals and to forswear grooming interactions with others, we can expect that grooming interactions among primates take the form of an optional Prisoner's Dilemma, where the opt-out payoff W is intermediate between R (the reward for mutual cooperation) and P (the penalty for mutual defection). To understand the evolution of mutual grooming under natural selection, we therefore need to consider the evolution of cooperative behavior in iterated optional Prisoner's Dilemmas.

Our illustration is loosely based on recent discussions of the social lives of primates (for example Smuts *et al.*, 1987). Plainly, it would be possible to test our assumptions about the ordering of costs and benefits by assessing the contributions to fitness of various commitments of time and of various states of parasitic infestation. Until such testing is done, we can only suggest that it is plausible that this part of primate social life may be understood in terms of optional Prisoner's Dilemmas. We also think it likely that optional games will prove useful in understanding cooperative hunting, cooperative foraging, mate-seeking coalitions, systems of defense against predation, and cooperation among females in exercising mate choice. But in all these cases the promise of the approach we recommend must be assessed in the light of field studies.

2. Analytic Results for Simple Strategies

Consider the following very simple strategies for playing an iterated optional Prisoner's Dilemma with the standard payoffs T , R , P , S and the opt out payoff W :

Solo: always opt out;

Undiscriminating Altruist: always interact and always play C;

Discriminating Altruist: interact with any animal

that is willing to interact with you provided that that animal has never previously defected on you, and cooperate in any such interaction, (or, if there are no such animals, opt out);

Undiscriminating Defector: always interact and always play D;

Discriminating Defector: interact with any animal that is willing to interact with you, provided that that animal has never previously defected on you, and defect in any such interaction (or, if there are no such animals, opt out).

We now consider, from a standard evolutionary game-theoretic perspective, the fitness relationships among some of these strategies under various conditions. Throughout we shall suppose that the population of animals with we are dealing has size N , and that the average number of occasions that an animal has to play the optional Prisoner's Dilemma during its lifetime is M . We assume that M is significantly larger than N , an assumption that we base on the idea that animals have frequent opportunities for playing the game and that they interact with a relatively small number of conspecifics. (Think of the number of times primates groom one another during their lifetimes and the relatively small sizes of the groups of conspecifics with which they interact.) Since our strategies involve the ability of animals to recognize one another, to recognize when another has defected on them, and to remember past defections, we are obviously supposing that our animals have substantial cognitive capacities. We are encouraged both by recent studies of the cognitive lives of primates (Cheney & Seyfarth, 1990) and by investigations of cooperative behavior in guppies (Dugatkin & Alfieri, 1991). However, these assumptions about the cognitive capacities of the animals concerned are plainly in tension with the "like begets like" assumption of evolutionary game theory (Maynard Smith, 1982). For the moment, we shall simply suppose that "like begets like" is a good approximate rule of thumb. Later sections will explore the dynamics of the evolutionary game in a less simplistic way.

RESULT 1. DISCRIMINATING DEFECTOR CAN BE INVADDED BY SOLO

Suppose that the population contains $N - n$ *Discriminating Defectors* and n *Solos* ($n > 0$). The payoff to each *Solo* is MW (on each of the M occasions on which the opportunity arises, a *Solo* opts out

for payoff W). The payoff to each *Discriminating Defector* is:

$$(N - n - 1)P + (M - N + n + 1)W,$$

since each *Discriminating Defector* plays with each other *Discriminating Defector* just once, each attempting to exploit the other and receiving payoff P ; when all have been tested, all opt out. Since $W > P$, the payoff to *Solos* is greater than that to *Discriminating Defectors*. A population fixed for *Discriminating Defector* should be invaded by *Solo*, and *Solo* should sweep to fixation.

RESULT 1.1. SOLO INVADES UNDISCRIMINATING DEFECTOR

The payoff to *Undiscriminating Defector* is less than that for *Discriminating Defector*. The payoff to *Solo* is the same as in Result 1.

RESULT 2. SOLO CAN BE INVADDED BY DISCRIMINATING ALTRUIST

In a population consisting entirely of *Solos*, a single mutant *Discriminating Altruist* is indistinguishable from the other members. If there are n (> 1) *Discriminating Altruists* in a population with $N - n$ *Solos*, then the payoff to each *Discriminating Altruist* is MR (provided that n is even) and $(n - 1)Mr/n + MW/n$ (when n is odd). (These provisions are needed since, with an odd number of *Discriminating Altruists*, if the decision opportunities arise for all simultaneously, arithmetical considerations dictate that one will have to opt out; the probability that any particular *Discriminating Altruist* is the unlucky one is $1/n$.) Given $n > 1$, it is trivial that the payoff to a *Discriminating Altruist* exceeds that to a *Solo*, for all values of n . Hence *Discriminating Altruist* can be expected to invade a population of *Solos*, and to sweep to fixation.

RESULT 2.1. WHEN BOTH DISCRIMINATING ALTRUISTS AND DISCRIMINATING DEFECTORS ENTER A POPULATION OF SOLOS, THE DISCRIMINATING ALTRUISTS MAY BE DRIVEN OUT; BUT IT IS ALSO POSSIBLE FOR THEM TO ENTER, GIVEN THE RIGHT ORDER OF MUTATIONS

If a population of *Solos* containing a single *Discriminating Altruist* mutant comes to have a single *Discriminating Defector* mutant, then the *Discriminating Altruist* mutant will be selected against. The condition for *Discriminating Altruists* to enter the population when *Discriminating Defector* mutations are also likely to arise is that the order of mutations be [*Discriminating Altruist*, *Discriminating Altruist*] rather than [*Discriminating Altruist*, *Discriminating Defector*]. If both mutations arise at the same frequency, we can expect that the probability that

Discriminating Altruists will drift into a population of *Solos* will be $p/2$, where p is the probability that *Discriminating Altruists* would increase in frequency in a population of *Solos* without *Discriminating Defector* mutations (i.e. a population like that described in Result 2).

RESULT 3. IF THE REWARD FOR COOPERATION IS LARGE IN RELATION TO THE PAYOFF FOR OPTING OUT, AND IF THE NUMBER OF DECISION OPPORTUNITIES IS LARGE IN RELATION TO POPULATION SIZE, DISCRIMINATING ALTRUISTS CAN INCREASE IN FREQUENCY AGAINST DISCRIMINATING DEFECTORS.

Suppose that a population contains n *Discriminating Altruists* and $N - n$ *Discriminating Defectors* where $n > 1$. The worst case for a *Discriminating Altruist* is to be exploited by each *Discriminating Defector* and then to spend the rest of the decision opportunities cooperating with other *Discriminating Altruists* (or, occasionally, opting out if the number of *Discriminating Altruists* is odd). This means that the payoff to a *Discriminating Altruist* is bounded below by

$$(N - n)S + (M - N + n)[(n - 1)R/n + W/n].$$

The best a *Discriminating Defector* can expect to do is to exploit each of the *Discriminating Altruists* once, and opt out on the remaining occasions. Hence the payoff to a *Discriminating Defector* is bounded above by

$$nT + (M - n)W.$$

The condition for *Discriminating Altruists* to increase in frequency under selection is thus

$$(N - n)S + (M - N + n)[(n - 1)R/n + W/n] > nT + (M - n)W.$$

This reduces to

$$M/N > K + Hn/N,$$

where

$$K = [(n - 1)R + W - nS]/(n - 1)(R - W);$$

$$H = [n(T + S) - (n + 1)W - (n - 1)R]/(n - 1)(R - W).$$

When n is large (approximately N), the crucial condition for the maintenance of *Discriminating Altruists* is

$$M/N > (T - W)/(R - W),$$

which is clearly satisfied if the reward for cooperating is significantly larger than the payoff for opting out, the payoff to exploiters is not too big, and the number of decision opportunities is sufficiently large relative to the population size. Hence it will be possible for *Discriminating Altruists* to resist invasion by *Discriminating Defectors*. At the other extreme, when n is small, the worst case is given by $n=3$. Since $N \gg 3$, we can ignore the term in H , and approximate the inequality by

$$M/N > (2R + W - 3S)/2(R - W).$$

As before, if the reward for cooperation is large in relation to the payoff for opting out, and if the number of decision opportunities is large in relation to population size, *Discriminating Altruists* can increase in frequency against *Discriminating Defector*.

These results are encouraging, for they suggest that, contrary to our naïve expectations, it might be very hard for anti-social or asocial behavior to be evolutionarily sustainable. By Result 1, anti-social populations are likely to decay into states of asociality. By Result 2, asocial populations are likely to be invaded by *Discriminating Altruists*, and, given Result 3, *Discriminating Altruists* can resist invasion by *Discriminating Defectors*. The only problem for a population of *Discriminating Altruists* is that *Undiscriminating Altruists* can drift in unnoticed, and, once the population has a sufficient number of them it is ripe for invasion by *Discriminating Defector* (or even by *Undiscriminating Defector*). Nevertheless, the combination of Results 1 and 2, and also Result 3, show that *Discriminating Altruists* can stage a comeback. Our analysis reveals that, while altruism may not be stable, the absence of altruism is also unstable. Moreover, when fitness differences are marked, we can expect that populations will spend most of their time in states of high cooperation, with occasional crashes and brief recovery periods. As we shall see later, this optimistic expectation is confirmed by computer simulations.

We now briefly explore the consequences of supposing that interactions are not fully optional. Animals are paired at random, and can either play with their assigned partner or opt out. As before, *Solos* always opt out, *Discriminating Altruists* play if and only if the assigned partner has not previously defected on them (and opt out otherwise) and they cooperate when they play, *Discriminating Defectors* also play if and only if the assigned partner has not defected on them and they defect when they play.

RESULT 4. DISCRIMINATING ALTRUISTS CAN INVADE A POPULATION FIXED FOR SOLO

Initially *Discriminating Altruists* are indistinguishable from *Solos*. Once there are two (or more) *Discriminating Altruists* there is a non-zero probability that they will be paired, and, on such occasions, each will receive R , a payoff that exceeds the opt out payoff W . So the fitness of *Discriminating Altruists* can be written as $M(rR + (1-r)W)$ where $r > 0$, which exceeds the payoff for *Solos* of MW .

RESULT 5. SOLO CAN INVADE A POPULATION FIXED FOR DISCRIMINATING DEFECTOR

In a population of *Discriminating Defectors*, the payoff to a *Discriminating Defector* will be $M(rP + (1-r)W)$, where $r > 0$ (r is now the probability that two *Discriminating Defectors* who have never previously met are paired). The payoff to *Solo* is MW , and, since $W > P$, *Solo* has a selective advantage.

RESULT 6. ALTHOUGH DISCRIMINATING ALTRUISTS CAN STILL INCREASE IN FREQUENCY AGAINST DISCRIMINATING DEFECTOR, THE CONDITION FOR DOING SO IN A SEMI-OPTIONAL GAME IS MORE STRINGENT THAN IN THE FULLY OPTIONAL GAME

Suppose the population has n *Discriminating Altruists* and $(N-n)$ *Discriminating Defectors*. The expected number of encounters of a *Discriminating Altruist* with a *Discriminating Defector* is $M(N-n)/N$. The total payoff to a maximally unlucky *Discriminating Altruist* from these encounters is $(N-n)S + M(N-n)W/N$. The rest of the time *Discriminating Altruists* are paired with one another for total payoff MnR/N . The total payoff for *Discriminating Altruists* is thus bounded below by

$$(N-n)S + M[(N-n)W + nR]/N.$$

The total payoff for *Discriminating Defector* is bounded above by

$$(N-n)T + (M - N + n)W.$$

Discriminating Altruists will have a selective advantage provided that

$$Mn(R - W)/N > (N-n)(T - S - W).$$

When *Discriminating Altruists* are prevalent (n is close to N), this condition becomes

$$M/N > (T - S - W)/(R - W),$$

which is very similar to the condition in the fully-optional game. When *Discriminating Altruists* are

rare ($n > 1$, but small in relation to N), increase of *Discriminating Altruists* requires that

$$2M/N^2 > (T - S - W)/(R - W).$$

This is far more exacting than the condition of Result 3, and it is thus only in special cases (M extremely large) that we could expect a small number of *Discriminating Altruists* to invade a population fixed for *Discriminating Defector* in the semi-optional game.

From Results 4–6, we can expect that the dynamics of the evolution of cooperative behavior in semi-optional games will resemble that in the fully-optional case, but that recovery from crashes is likely to be slower and mediated by the presence of *Solos*. Intuitively, the direct route from *Discriminating Defector* to *Discriminating Altruist* is now partially blocked, and, very frequently, only the trajectory from *Discriminating Defector* to *Solo* to *Discriminating Altruist* will be available.

3. Computational Simulations

The above analytical results are based on assuming that the populations are relatively simple in two ways: first, that the populations consist of individuals who use one of a small set of strategies; second, that the strategies are chosen from the set of simple strategies described in Section 2. In order to investigate the properties of more heterogeneous populations, and of populations containing individuals following more complex strategies, we performed a number of computational simulations of populations of players who participate in the iterated Prisoner's Dilemma by following inherited strategies. Our simulation results support and expand upon the analytic results, and also illustrate how the genetic representation of strategies can influence the evolutionary dynamics of populations whose members deploy those strategies.

This section describes the algorithms used in the computational simulations. The next section presents the results of those simulations, and describes the dynamics of a few of the runs we performed.

In the simulations, the actions performed by each player are represented as a history sequence. The lengths of histories recorded in our simulations varied from 2 to 4. For example the sequence (C C) indicates that a player cooperated on both of the previous two rounds; the sequence (D C) indicates that a player defected two rounds ago, but cooperated the last round. The symbol N is used when the players haven't played as many rounds as the history records. Thus for a history of length 2, the sequence (N N) indicates that no rounds at all have been played; the sequence (N C)

means that a single round was played, and the player cooperated.

Strategies are represented by pairing each possible history of opponent's actions with the action to make in response the next round. This pairing of a history with a response action will be called a "move". The following move represents the response of defecting if the opponent cooperated twice in a row:

((C C) D)

This move represents the action of cooperating in the first round:

((N N) C)

Given a specific history length, a complete strategy contains a move for each possible history sequence of that length. For example this strategy represents the *Tit-for-Tat* strategy in which two steps of history are recorded:

((N N) C)
((N C) C)
((N D) D)
((C C) C)
((C D) D)
((D C) C)
((D D) D)

In this strategy, the player begins by cooperating and then responds with whatever its opponent did the last round.

A sequence of rounds between two players is simulated by using the strategies of the two players to determine their moves for each round, depending on what the other player did the last rounds. Each player receives an increment to a "fitness" value according to this payoff schedule:

Payoff	Explanation	Value
T	Defect if other cooperates	7
R	Both cooperate	5
W	Opt out	3
P	Both defect	2
S	Cooperate if other defects	0

In a generation, each player plays against each other player in the population some number of rounds.

For example consider a simulation of the compulsory game in which one step of history is recorded. Two players with the following strategies are chosen to play the game:

((N) C)	((N) D)
((C) C)	((C) D)
((D) D)	((D) C)
Player one	Player two

The players begin their interactions with fitness values of 0. In the first round both of the histories are (N), so player one will cooperate and player two will defect. Player one's fitness will remain 0, while player two's fitness will be set to 7. In the second round, player one consults its strategy with the history (D) and defects. Player two uses the history (C) and also defects. Both players fitness values are incremented by 2. On the third round, player one defects again, but player two cooperates. Hence player one receives 7 and player two receives 0. The next round player one cooperates again and so does player two, so they both receive 5 points. The next round player one cooperates but player two defects. After this point the histories are identical to that after the first move, and so the players perform as they did then, and continue to cycle through the same sequence of moves.

In the first generation of a simulation, each of the players is assigned a random strategy—the response action to each possible history sequence is randomly chosen from the available moves. At the end of each subsequent generation, the set of players is sorted in order of decreasing value of the total fitness payoffs each received while playing against the other players. The top third of the players is preserved into the next generation, and those players are also used to create the strategies of the rest of the players in the next generation. Each new strategy is created by mixing the strategies of two of the most successful players—for each possible history, the response action is taken randomly from one or the other parent's strategy. Mixing strategies in this way has the effect of rapidly distributing advantageous moves through the population. A small fraction (for most of our runs: 1%) of the moves are then mutated by replacing the action part of the move with a randomly chosen action.

In each generation, a record is kept of the total number of moves of each type: “cooperate”, “defect”, and for the optional games, “opt out”. At the end of a generation the average fitness of the population is also recorded. A sample run of the compulsory game is shown in Fig. 1. As is typical for the runs reported here, the population moves through a number of states in which the levels of cooperation and defection are fairly stable for tens of generations or longer. Rapid transitions then occur, yielding other stable states.

In this run the population quickly enters a state of high cooperation and fitness. Around generation 70, it reverts to a state of virtually 100% defection and low fitness. This is followed at generation 150 with a state of 50% cooperation, 50% defection and an intermediate fitness value. Around generation 310, the population again finds a state of very high cooperation.

The results of an entire run are summarized with two numerical values:

The “cooperativity” measure is meant to quantify the degree to which cooperative behavior dominated during the run. Cooperativity is defined as the fraction of generations during a run when the difference between the percentages of cooperative moves and defection moves is greater than a threshold of 25. The cooperativity value for the run shown in Fig. 1 is 0.494. (The precise value of the threshold for computing the cooperativity value is not crucial. For example changing the threshold to 70 for the run in Fig. 1 changes the cooperativity value from 0.494 to 0.488. This is because the runs tend to remain in states where either cooperation or defection is relatively high, and the other is correspondingly low.)

The “instability” measure is meant to quantify the degree to which the amount of cooperation varies from generation to generation. This is defined as the average of the square of the difference between the number of cooperative moves in successive generations. The instability value for the run shown Fig. 1 is 19.41. The instability measure increases if a run enters more states, or if the states that it enters do not have constant values of cooperation. The value of the instability measure for simulations depends on the simulation parameters. For example in a set of 20 runs of 500 generations of the compulsory game with 36 players, a history of length 2, and 10 rounds between each pair of players, increasing the fraction of moves mutated each generation from 0.1% to 10% changed the average instability value from 3.28 to 28.37. As is shown below, the instability value is also strongly affected by whether the game is compulsory or optional.

Two versions of the optional game were simulated: the “semi-optional” and the “fully-optional” game. In the semi-optional game players are paired off as in the compulsory game, and each pair plays some number of rounds against each other. If either player chooses to opt out in any move, both players receive the opt-out payoff W .

In the fully optional game, players first attempt to locate other players who will not opt out against them. Such pairs then play one round against each other as in the compulsory game, and each player keeps a record of what its opponent does which it consults the next time they are looking for partners. All players who do not find a willing partner in a given round are awarded the opt-out payoff. The fully optional game requires more computational overhead to simulate, and the runs take much longer, than the semi-optional game, because each player must record all of the histories of its interactions with all of the players it has played against in a generation; and because the process

TABLE 1

Game	Mean cooperativity	Standard deviation	Mean instability	Standard deviation
Compulsory	0.105	0.160	5.28	3.80
Semi-optional	0.719	0.283	13.2	5.96

of pairing off the players is more complicated than for the semi-optional game.

Runs of the optional games are illustrated and discussed below.

4. Simulation Results

The results of this study can be summarized in Tables 1 and 2. The statistics in Table 1 are for a set of 27 runs, each of 500 generations. There were 60 players in each run, two steps of history were recorded, and each pair of players played 30 rounds each generation. As can be seen from Table 1, the compulsory game yields populations which are more stable but less cooperative than those that play the semi-optional game.

Fewer runs of the fully optional game were done, as they take much longer. However the general result is similar. Table 2 shows results for 12 runs of each of the indicated games of 500 generations each, with 30 players, two steps of history were recorded and the players played against each other an average of ten rounds.

The reason that the statistics for the compulsory game in Table 2 are different from those in Table 1 is that the algorithm for pairing off players in this set of runs corresponded to that used for the fully optional game, except that no player could opt out. In each round of the game, a player was paired with another player in the population at random. The fact that some pairs played fewer (and some played more) rounds than the average of ten is reflected in the higher instability values for these runs as compared with the first table. Again however, the optional version of the game yields more cooperative generations than does the compulsory version of the game, and the instability of the optional game is higher than that of the compulsory game.

4.1. RUNS OF THE COMPULSORY GAME

Runs of the compulsory game tend to become stuck in a small number of states, either with very high cooperation, very high defection, or half cooperation and half defection. Often a run will be stuck in a state for many generations. This is reflected in the relatively low instability value for the compulsory game, and the high ratio of the standard deviation of its cooperativity to the mean value.

For example, when one step of history is recorded, a run often first enters a state where all of the players defect each round. This is because the initial strategies are random, and so there are a large number of players who cooperate no matter what their opponents do. Hence the defectors receive the high T payoff, and their offspring take over the population.

Subsequent events can be understood by examining the strategies shown in Fig. 2, beginning with the *Undiscriminating Defector* strategy shown in Fig. 2a. Two of the single-move mutations of this strategy will be at a disadvantage playing against it because they will cooperate in a round where the original will defect. If everyone in the population is defecting all the time, however, the move ((C) D) is never exercised. So a mutation from the move ((C) D) to ((C) C) will not affect the behavior of (nor the fitness payoffs accumulated by) a player using the strategy.

After several generations of this kind of “genetic drift”, a population initially containing only players with strategy 2a can be expected to contain a fraction of players with strategy 2b. From here, a single mutation in the ((C) D) move can change the strategy to *Tit-for-Tat*, as shown in Fig. 2c. Provided that enough of these mutations occur at about the same time, players using *Tit-for-Tat* can dominate the population, which will enter a state of very high cooperation. This is essentially what happens in the run shown in Fig. 1 around generation 10.

TABLE 2

Game	Mean cooperativity	Standard deviation	Mean instability	Standard deviation
Compulsory	0.248	0.267	14.8	8.85
Fully-optional	0.668	0.292	29.9	5.10




FIG. 1. A run of the compulsory game, with two steps of history recorded. The trace marked C records the percentage of "cooperate" moves each generation; the trace marked D records the percentage of "defect" moves each generation; the trace marked F records the average fitness of the population as a percentage of the maximum possible value. The "cooperativity" value for this run is 0.494; its "instability" is 19.41.




FIG. 4. A run of the semi-optional game, with one step of history recorded. The trace marked O records the percentage of "opt out" moves each generation. The "cooperativity" value for this run is 0.581; its "instability" is 45.1.




FIG. 5. A portion of a run of the semi-optional game, showing two "predator-prey" cycles, and the beginning of a state of high cooperation.




FIG. 6. A run of the fully optional game with two steps of history recorded. The "cooperativity" value for this run is 0.716; its "instability" is 31.7.

((N) D)	((N) D)	((N) C)	((N) C)	((N) C)	((N) D)
((C) D)	((C) C)	((C) C)	((C) C)	((C) D)	((C) D)
((D) D)	((D) D)	((D) D)	((D) C)	((D) C)	((D) C)
a	b	c	d	e	f

FIG. 2. Some strategies for the compulsory game, when one step of history is recorded.

However *Tit-for-Tat* is not immune to variation. For example the ((N) C) move could mutate back to ((N) D) and return the population to strategies like 2b and a state of high defection. This is what happens near generation 70 in the run shown in Fig. 1.

Another mutation from *Tit-for-Tat* can yield the *Undiscriminating Altruist* strategy 2d. If the population is in a state where everyone else is following *Tit-for-Tat*, following the *Undiscriminating Altruist* strategy will not affect the fitness of the player, as the ((D) C) move in *Tit-for-Tat* is never exercised. However if defectors appear, they can exploit the player who always cooperates. For example, strategy 2e is one mutation away from *Undiscriminating Altruist*. If it appears, the number of players using it will increase in the population as they prey on the cooperators. A single mutation from 2e is 2f, which increases defection in the population even more.

Strategies like the one shown in Fig. 2f can lead to very stable populations, with relatively low cooperation and fitness values. A pair of players playing this strategy will cooperate 50% of the time and defect 50% of the time. Furthermore there is no possibility of genetic drift with this strategy as each move of it is exercised and each one-mutation variant of this strategy is at a disadvantage against it. The population in the run shown in Fig. 1 enters a state where each member of the population is playing a variant of this strategy around generation 155. Note that the state remains steady for around 60 generations, before the population returns to a state of relatively high cooperation.

With longer history lengths, strategies like 2f, in which pairs of players using the strategies alternate between cooperation and defection, can be very stable. For example the strategy shown in Fig. 3 is a variant of 2e for a history length of 2. In Figure 3, the '*' character indicates moves that are exercised when players using this strategy play against each other. Five of the seven moves are used in such an encounter, and therefore cannot mutate away from those shown without having an immediate effect. Mutations to the other two moves actually serve to reinforce the patterns of interactions seen when the rest of the players are using this strategy. It is important to note that the stability of this strategy is not a matter of any static superiority of the strategy compared with all possible competitors: instead this strategy is stable

because its genetic representation is such that there are no better alternatives to it via single mutations. So a population in which all players are following such a strategy constitutes a robust local optimum in the evolutionary search. We have observed runs of the compulsory game in which the population remains stuck in states where populations are using such strategies for thousands of generations, though transitions to states of high cooperation or high defection eventually occur.

4.2. RUNS OF THE OPTIONAL GAMES

With the option of opting out, populations in the optional games can escape from states in which there is a significant amount of defection. Since the payoff for opting out, W , is larger than both S , the payoff for cooperating when the opponent defects, and P , the payoff for mutual defection, the presence of defection in the population makes opting out an advantageous alternative. Thus populations playing the optional games will tend to revert to states of high opting out whenever a number of defectors appear. This fact alone accounts for some of the reason why the optional game leads to higher cooperativity—it just can't become stuck in states of high defection.

A run of the semi-optional game with one step of history recorded is shown in Fig. 4. As is typical for runs of the optional games, the population enters more states than otherwise equivalent runs of the compulsory game, and the states that it enters are much less stable.

Since the initial strategies are random, and therefore include many defectors, opting out is initially favored, and most runs of the semi-optional game enter states of virtually 100% opting out in the early generations, often as early as generation 10. From then one, the populations tend to go through cycles of various sorts. One kind of cycle involves the appearance of defectors,

((N N) D)	*
((N C) C)	
((N D) C)	*
((C C) D)	*
((C D) C)	*
((D C) C)	*
((D D) D)	

FIG. 3. A stable sub-optimal strategy for the compulsory game, when two steps of history are recorded. The moves marked * are exercised when a pair of players with this strategy compete against each other.

TABLE 3

Optional game	Mean cooperativity	Standard deviation	Mean instability	Standard deviation
Semi-	0.401	0.263	27.8	10.02
Fully-	0.668	0.292	29.9	5.10

which is usually soon followed by a reversion to high opting out.

Another kind of cycle seen in the semi-optional games involves the appearance and subsequent disappearance of players who always cooperate. If almost every other member of the population is opting out each round, there is no danger to the few players who mutate to strategies that involve some cooperation. Indeed, when these players play against each other, and receive the reward R for mutual cooperation (which is larger than the W opting out payoff), they will increase in the population. However when there are lots of careless cooperators in the population, there is an advantage to be gained by defecting. If mutations occur to create strategies that involve some defection, defectors will rapidly increase, effectively destroying the cooperators. Since the resultant level of defection is high, opting out is now relatively advantageous, and the population reverts to a state where everyone is opting out. This pattern is similar to the “predator–prey” cycles seen in population biology.

A typical run of the semi-optional game will go through a number of cycles. In some cases it is possible that the mutations that increase defection either include or are followed by mutations that increase the discriminatingness of the strategy, either by playing *Tit-for-Tat*, or by opting out when an opponent defects, i.e. the *Discriminating Altruist* strategy. If such mutations occur before defection rises significantly, it is possible for the players possessing these strategies to continue to increase in the population even when defection rises temporarily. Thus the population can enter and remain in a relatively stable state of high cooperation.

This process is illustrated in Fig. 5. This shows a portion of a run of the semi-optional game. At generation 60, virtually all of the players are opting out in each game. Around generation 70, a few players begin cooperating. Since they manage to find other cooperators, their numbers increase. Within two generations, however, a few defectors appear. Since these defectors will prevail over the cooperators, their numbers increase rapidly, and by generation 80 or so, the cooperators are gone. A similar but less dramatic pattern of this sort begins almost immediately, and is over by generation 90.

At generation 91, another set of cooperators appears, followed closely by defectors. However in this case, at least some of the cooperators are playing a discriminating strategy, and in fact by generation 108, the defectors begin disappearing from the population. By generation 110, virtually all of the players are cooperating all of the time.

As with the compulsory game, states of high cooperation are not stable either. With all of the members of the population cooperating, genetic drift can set in, and mutate some of the discriminating strategies to their careless versions, providing fodder for defectors when they appear by mutation. As before, the high rate of defection will ultimately be followed by an increase in opting out.

The general dynamics of the fully optional game are similar to those for the semi-optional game. Slightly higher cooperativity is seen, as illustrated in Table 3, but it is not clear if this is significant. The parameters for these runs are: 30 players, 10 game interactions, history length of 2. The statistics are for 12 runs of 500 generations each.

A run of the fully optional game with 2 steps of history recorded is shown in Fig. 6. The population quickly finds a state of high cooperation without an intervening period of opting out, as predicted by the analysis in Section 2. Around generation 70, this state crashes and yields a period of high opting out that lasts (with one short glitch) until generation 205. At this point, the population enters a state where cooperation is still relatively high, but the fraction of defection moves is around 30%. This state lasts until generation 440, when a state of high cooperation occurs. Again, the transition to the higher level of cooperation happens without an intervening period of opting out.

5. Conclusion

The superiority of the optional games in reaching states of high cooperation can be demonstrated analytically, and is supported by the dynamic properties that simulations of such games manifest. There is no way for a population playing the compulsory game to escape from a state of high defection, except if several favorable mutations appear simultaneously. In the optional games there are routes

out of states of high defection. The option of asocial behavior facilitates the appearance and maintenance of altruistic behavior.

In thinking about the evolution of social behavior it is important to recognize that such behavior occurs against a changing environment consisting of the behaviors of the other members of the populations. Thus such an evolutionary process is a feedback system, and the global properties of such a process should be expected to fluctuate, perhaps chaotically. The relative fitness of a given behavior or strategy cannot be assessed statically, with respect to a specific, or to a fixed, environment. In the long run, the evolutionary dynamical properties of strategies and their genetic representations, may have the most significant effect on the careers of populations using those strategies.

Obviously, in addition to more detailed analysis and simulation of the optional games, it is important to see if the optional games provide a more ethologically valid model of some animal interactions. One would have to be able to distinguish between an animal's refusal to participate in an interaction ("opting out") and its failure to reciprocate an altruistic action of another animal ("defection"). Our model predicts that geographically separate, but genetically equivalent, populations of the same species might differ markedly in their social interactions, with some populations exhibiting high cooperation, some behaving asocially, and others enduring periods of highly antisocial behavior.

There are many ways in which to introduce complications into the study of altruism in optional and semi-optional games. For example, animals may make various types of errors of recognition, potential partners may vary in quality, and different types of game-theoretic situations may arise. Our preliminary analyses reveal that these complications do not markedly affect the results presented herein. A slightly more detailed survey of some possible complications is given in Kitcher (1993), where the evolution of specifically human types of altruism is also addressed.

REFERENCES

- AXELROD, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- AXELROD, R. & HAMILTON, W. D. (1981). The evolution of cooperation. *Science* **211**, 1390–1396.
- BOYD, R. & LORBERBAUM, J. P. (1987). No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game. *Nature* **327**, 58–59.
- BULL, J. J. & RICE, W. R. (1991). Distinguishing mechanisms for the evolution of co-operation. *J. theor. Biol.* **149**, 63–74.
- CHENEY, D. L. & SEYFARTH, R. M. (1990). *How Monkeys See the World*. Chicago: University of Chicago Press.
- DUGATKIN, L. A. & ALFIERI, M. (1992). Guppies and the tit for tat strategy: Preference based on past interaction. *Behav. Ecol. Sociobiol.* **28**, 243–246.
- DUGATKIN, L. A. & WILSON, D. S. (1991). Rover: A strategy for exploiting cooperators in a patchy environment. *Am. Nat.* **138**(3), 687–701.
- FARRELL, J. & WARE, R. (1989). Evolutionary stability in the repeated prisoner's dilemma. *Theor. Popul. Biol.* **36**, 161–168.
- KITCHER, P. S. (1993). The evolution of human altruism. *J. Phil.* **90**, 497–516.
- LIMA, S. L. (1989). Iterated prisoner's dilemma: an approach to evolutionarily stable cooperation. *Am. Nat.* **134**, 828–834.
- LINDGREN, K. (1992). Evolutionary phenomena in simple dynamics. In: *Artificial Life II* (Langton, C. G., Taylor, C., Farmer, J. D. & Rasmussen, S., eds) pp. 295–311. Redwood City, CA: Addison-Wesley.
- MAYNARD SMITH, J. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- MAYNARD SMITH, J. & PRICE, G. R. (1973). The logic of animal conflict. *Nature* **246**, pp. 15–18.
- MESTERTON-GIBBONS, M. (1992). On the iterated prisoner's dilemma. *Bull. math. Biol.* **54**, 423–443.
- MILLER, R. R. (1967). No play: a means of conflict resolution. *J. Pers. & Soc. Psych.* **6**(2), 150–156.
- NÖE, R., VAN SCHAİK, C. P. & VAN HOOF, J. (1991). The market effect: an explanation for pay-off asymmetries among collaborating animals. *Ethology* **87**, 97–118.
- PECK, J. R. (1990). Evolution of outsider exclusion. *J. Theor. Biol.* **148**, 563–571.
- PECK, J. R. & FELDMAN, M. W. (1986). The evolution of helping behavior in large, randomly mixed populations. *Am. Nat.* **127**, 209–221.
- SMUTS, B. B., CHENEY, D. L., SEYFARTH, R. M., WRANGHAN, R. W. & STRUHSACKER, T. T. (eds) (1987). *Primate Societies*. Chicago: University of Chicago Press.
- STANLEY, E. A., ASHLOCK, D. & TESFATSION, L. (1994). Iterated prisoner's dilemma with choice and refusal of partners. In: *Artificial Life III* (Langton, C. G., ed.) pp. 131–175. Redwood City, CA: Addison-Wesley.
- TRIVERS, R. (1971). The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 33–57.