

The Prisoner's Dilemma and Predator-Prey Coevolution

Joel Pomerantz
Willard L. Miranker
Yale University
Department of Computer Science
TR-1214

The Prisoner's Dilemma and Predator-Prey Coevolution

Joel Pomerantz, Willard Miranker
Department of Computer Science, Yale University
joel.pomerantz@yale.edu, miranker@cs.yale.edu
New Haven, CT 06511

Abstract

Iterated Prisoner's Dilemma is a game designed to model open-ended interactions involving competition and cooperation. However, with the introduction of faulty information, players become both more aggressive and more forgiving. This allows the game to model exploitation as well. In this paper, the use of faulty information, spatialization, genetic algorithms, (including a novel modeling of mitochondrial genes which influences predator-prey and speciation issues) and several reproductive strategies are explored in terms of their use in modeling predator-prey interactions and their coevolution. Interactions between the fungal and algal components of lichen are explored, as an application.

Keywords: iterated prisoner's dilemma, genetic algorithms, mitochondrial genes, spatialization, speciation, noise, coevolution

1. INTRODUCTION

Iterated Prisoner's Dilemma has been explored in detail under conditions of perfect information but has been less extensively studied with noise. Noise adds an element of uncertainty to the game, which makes the interaction of strategies nondeterministic, and which makes so-called exploitative strategies slightly more feasible. Additionally, the combination of noise and spatialization (the placing of players in a tournament into slots in a grid where proximity influences the likelihood of interaction) is novel and illuminating. This is because the use of spatial grids in the Prisoner's Dilemma is typically employed to save small amounts of processing time, whereas the introduction of noise requires additional amounts of processing power next to which the savings from spatialization are negligible (Martino, 1995). However, the joint use of noise and spatialization allows study of local responses of strategies to noise in relationship to neighbors' responses. This is useful because it allows small groups of individuals to interact without having to also interact with all the other individuals in play in the tournament. Because the strategies are generated by genetic algorithms, they are able to continually evolve in response to the strategies proximal to them in the grid, and thus to be more responsive to their neighbors play (Axelrod, 1987). The conjunction of noise, spatialization, and genetic algorithms thus allows strategies to evolve based on their neighbors rather than the set of all strategies in existence, to develop means of exploiting weaknesses in these neighbors (i.e. local conditions), as well as to respond to the exploitative practices of their neighbors. Together, these three features provide an interesting framework for modeling predator-prey coevolution. The evolution of altruism (an individual acting in a way that helps its neighbors at its own expense, potentially helping to propagate the "altruistic" gene) is examined, as are the conditions

under which interactions between members of distinct species tend to be symmetric (a quality relevant to our modeling of symbiosis). Our genetic algorithmic approach makes use of two classes of genes; the customary so-called nuclear genes and a novel class modeled on mitochondrial genes. The latter inform both predator-prey issues and speciation in our Prisoner's Dilemma model. To illustrate our approach, interactions between fungal and algal components of lichen are modeled as a Prisoner's Dilemma tournament.

2. HISTORY

Flood and Dresher invented the Prisoner's Dilemma in 1950, as a model for cooperation. They originally framed it as a paradox. Suppose two prisoners could not communicate with one another, and were being asked to confess to a crime. Should only one confess, he would be let off while his partner in crime would receive a rather heavy sentence. Were both to confess, both would receive a moderate sentence. If neither confessed, both would get a light sentence. The paradox arises in that it is in either player's interests to confess no matter his partner's decision, so both will confess. Yet, both players would be better off had they both remained cooperated and remained silent rather than confessing.

The Prisoner's Dilemma has since been adapted by numerous authors as a model for the difficulty of securing cooperation when there exists an incentive to defect (corresponding to the prisoner confessing). In most adaptations [Axelrod, 1987; Sandholm, Crites; Martino, 1995), the players are no longer prisoners, and generally play to maximize positive consequences rather than minimize negative ones. For example, Axelrod sets up a payoff matrix of 3 points for mutual cooperation, 1 point for mutual defection, 0 points for a sucker (a cooperator whose opponent defects), and a 5 point "temptation to defect" ("betrayal" or unilateral defection) payoff (Axelrod, 1984, p. 8). The same dynamic is preserved: it is in either player's interests to defect regardless of the other player's choice, but both would be better off cooperating.

3. DEFINITIONS, TERMINOLOGY, COMMENTARY

3.1 Parameters

The paradox is changed into a more interesting model/game by iterating it. In many circumstances, two actors have multiple opportunities to interact with one another. Politicians must interact on several bills, corporations conduct a sequence of business deals, and successive generations of microorganisms will encounter one another. In such circumstances, cooperative arrangements may arise. This is not, as one might think, because greater familiarity means one can trust another to cooperate. Rather, it is because one's own actions will affect the other's subsequent actions. When the game is played for one round, Player A's decision can not affect Player B's decision. When the game lasts several rounds, it might. Depending on the strategy player B is using, player A might find that cooperating now will gain for himself greater rewards in the long run than defecting. This depends, of course, on player B's strategy and the likelihood the game actually lasts several rounds.

3.2 Rounds/turns, games, and tournaments

A tournament is a simulation of overall behavior of strategies, where each player plays several games against one or more opponents. A strategy is an algorithm that produces one decision in each round (turn) of a game (a decision is to cooperate with probability p , and to defect with probability $1-p$), given all the information available to it. Since confusion should

not occur, we shall use the terms "a strategy" and "a player employing a strategy" interchangeably. Frequently, the strategies may evolve between games, although this is not always true. Each game consists of one or more rounds of play; during each round of the game the two strategies decide independently whether to cooperate or defect. Behavior in one round impacts behavior in future rounds by giving the strategy additional information; behavior in one game can (in certain types of tournaments) affect behavior in future games by changing the strategies of the players involved. It is reasonable not to change a strategy during a round.

It is interesting to note, as Axelrod demonstrates, that the Prisoner's Dilemma is widely useful over a large variety of parameters. For example, the players only need to be assumed to care about their own interest, the payoff matrix need not be symmetric for both players or even comparable, and communication between players need not be prevented, as long as both players value a successful betrayal most, then a mutual cooperation, then a mutual defection, and least a sucker's payoff.

Axelrod (Axelrod, 1997) describes a bewildering array of uses for the Prisoner's Dilemma, from arms races, oligopolistic competition, and vote trading, to aggression in individuals, to westernization in Central Africa, etc. For every given model, the appropriate parameters of the game are different. Payoffs, lengths of games, amount of information shared, and initial starting conditions are all highly dependent on the situation for which the Prisoner's Dilemma is being used. There are some broad claims that can be made about the majority of tournaments, which apply over a range of situations. However, some of the most important findings necessarily depend on the specific parameters of the tournament.

One universal finding is that the number of rounds remaining in a game is very important. In single-round games, the optimal strategy is ALWAYS DEFECT. (Here and hereafter, the name of a strategy will be denoted in upper case, and the meaning of that strategy will usually be clear from its name.) In games where decisions are iterated an indeterminate number of times, cooperation becomes a much better option. If in each game, a player bases its decisions on its opponent's decisions, the opponent has more reason to cooperate. After all, if one knows one's opponent is using the so-called TIT FOR TAT strategy, one's own best move is to cooperate as long as the amount gained in the round of defection is not expected to exceed the amount lost due to the opponent's defection in the following round (Axelrod, 1984).

As the number of rounds remaining decreases, cooperation becomes less profitable. After all, at the end of a series of interactions, short-term gains become large relative to the expected rewards of continued cooperation. This tends to be seen in many political and biological examples. Savage (Savage, 1977) shows harmless bacteria tend to reproduce wildly in response to certain illnesses, which further harms the host (and thus harms chances of long-term cooperation), but increases the short-term chances of the bacteria spreading to other hosts.

A very interesting example of how longer-term interaction leads to cooperation is found in World War I. When facing enemy soldiers, one can fire to kill, or fire harmlessly. Normally, it is in one's interests to fire to kill, as a weaker enemy line is less dangerous to one's own wellbeing. Thus, in most wars one sees soldiers firing at one another in earnest. In the trenches of World War I, however, the same soldiers faced one another day after day. The commanders insisted that the soldiers shell one another, but in many places an arrangement of "live and let live" prevailed. Hay (Hay, 1916) describes a German force, quoting Thomas Atkin's diary which characterizes "their offensive operations with a tactful blend of constant firing and bad shooting, which while it satisfies the Prussians causes no

serious inconvenience to Thomas Atkins". Each side knows that to shoot to harm would provoke retaliation, and thus refrains from such behavior. Interestingly, as the war progressed and units were rotated more frequently, the fragile arrangements of cooperation tended to break down. When one does not face the same opponents for prolonged periods, cooperation is less profitable and therefore less likely.

3.3 TIT FOR TAT

The TIT FOR TAT strategy initially cooperates, and then on each round makes the move its opponent just used. Thus, two TIT FOR TAT players will always cooperate with one another; a TIT FOR TAT faced with a strategy that always defects will initially cooperate, and then defect on all subsequent rounds. TIT FOR TAT thus can never get more points than its opponent. For it to get "successful betrayer" points, the opponent would have had to begin a cycle of defections, which means the opponent would also have received these betrayer points. Nevertheless, a large variety of models for cooperation using the Iterated Prisoner's Dilemma (Axelrod, 1997) demonstrate that the strategy of TIT FOR TAT is usually the best strategy. It does very well against another TIT FOR TAT or other so-called cooperative strategies, and also does nearly optimally against strategies that frequently defect. Moreover, TIT FOR TAT tends to win tournaments without winning games (Axelrod, 1984).

In a wide variety of experiments, nice strategies usually win. A nice strategy (Axelrod, 1984, pp. 33-36) is one that does not defect until the other player has defected. However, niceness tells one nothing about forgiveness (willingness to return to cooperation after one's opponent defects). A strategy which will always defect once its opponent has defected once is still nice, although this one has not historically done particularly well in most tournaments due to the presence of strategies which "probe" their opponent to see if they can get away with defection, and then cooperate if they find they cannot (Axelrod, 1984, p. 37).

It is very important to note that the presence or absence of certain opponent strategies is what makes one strategy better than another. In a world where a strategy will be playing primarily against opponents who choose their actions randomly, ALWAYS DEFECT becomes the best strategy. What is interesting about TIT FOR TAT is that it does rather well relative to its neighbors in a large variety of environments, whether those neighbors are primarily nice or non-nice. However, it is frequently the case in Prisoner's Dilemma tournaments that TIT FOR TAT wins primarily due to the presence of "king-maker" strategies (Axelrod, 1984, p. 34): strategies which may do poorly themselves, but influence the determination of the eventual winner.

3.4 Spatialization

Spatialization is a tool that allows strategies to interact locally, moreover within the context of a larger environment. In the majority of tournaments, every strategy plays every other strategy. This means that to be successful a strategy must do well against every other strategy present. The implementation used here, however, sets up strategies with locations in a two-dimensional grid (sometimes termed, the board). The chance any two strategies will play one another in a particular game is reduced as distance between the strategies increases. This means that strategies must do well playing their neighbors in the grid, but need not be competitive with every existing strategy. Of course this is a feature resembling the struggles in nature.

3.5 Genetic algorithms

Genetic algorithms are a means of solving a variety of mathematical problems, particularly function minimization. They employ a coded sequence, termed a genetic sequence, which represents an approach to the problem in question. For us the sequence will encode a strategy. Each sequence is tested according to some test function¹, and those that do poorly² are eliminated. The remaining sequences are permitted to reproduce (i.e. to create new copies of themselves) with introduced mutations (small changes in the coded sequence), so that the solutions tested thereafter are similar but not numerically identical to the solutions which previously did well and thus survived elimination. The use of genetic algorithms means that the strategies found in abundance in a tournament are those which did well previously. If ALWAYS COOPERATE does very poorly, then elimination leaves fewer ALWAYS COOPERATE strategies to be faced later in the tournament. If TIT FOR TAT does well, then more and more TIT FOR TAT players will be found. This means that ALWAYS DEFECT might do quite well initially if there are many ALWAYS COOPERATE players around, but by eliminating these latter players, it will begin to do more poorly.

After each round in a genetic tournament, a number of players are killed (removed from play). Each player's chance of dying is related to the death rate of the game (a universal feature determining how likely players are to die), as well as its fitness rating (altered cumulatively over the course of the game by interactions with other players). Each player also may reproduce, generating an imperfect (mutated) copy of itself. In sexual reproduction, this "copy" is an imperfect blend of the coded sequences (the strategies) of two players. Because the tournaments described here are spatialized, reproduction takes place only into empty slots in the board (the grid). In such tournaments, an empty board slot may be chosen as the location of a new strategy; the likelihood of a player being a parent of the new strategy is related to its fitness and to its proximity to the empty slot.

Note that spatialization makes it easier for certain strategies to increase in number despite performing poorly against the majority of strategies present. For example, in a world (set of strategies likely to be encountered) dominated (populated for the most part) by ALWAYS DEFECT, a small number of TIT FOR TAT will not do particularly well. They must do slightly worse (i.e. acquire a lower fitness score) than the ALWAYS DEFECT whenever they face an ALWAYS DEFECT, and much better whenever they face a TIT FOR TAT. It can be shown that when the ratio of TIT FOR TAT to ALWAYS DEFECT is low, the TIT FOR TAT will never take hold (i.e. enlarge in number). However, if a local concentration of TIT FOR TAT is high under spatialization, this will permit those TIT FOR TAT to do well. A small cluster of the new strategy will form within the larger region of ALWAYS DEFECT (Axelrod, 1997), and can then expand.

Similar features are also found in so-called ecological tournaments, where a set of initial strategies play one another and then reproduce based on their success. Their difference from genetic algorithms is that reproduction is asexual and does not involve mutation, and so players in an ecological tournament generate exact copies of themselves during reproduction. This is important, because genetic algorithms allow new strategies to arise throughout a tournament. To win an ordinary ecological tournament, a strategy must

¹ The test function for an approach may or may not be based on the original problem. For Prisoner's Dilemma tournaments, the test is to actually play a game of Prisoner's Dilemma with another strategy.

² The means of evaluating the quality of an approach can vary widely. For these Prisoner's Dilemma tournaments, a player has a "fitness" score, which is cumulatively changed according to the points it gains or loses by cooperation or defection in the games played.

do sufficiently well during every time period to at least survive, and then extremely well near the end of the game. Using genetic algorithms, a strategy might be wiped out early on, but if it is the best strategy later in the game, it can still spontaneously reappear and then spread.

Thus, for Iterated Prisoner's Dilemma tournaments, genetic algorithms offer two significant advantages. First, a wider range of strategies can arise and be tested than a person can realistically devise or store subject to space limitations. A tremendous number of possible strategies can be encoded by any reasonably long code, and to test each of these separately would be an enormous task. Second, strategies can evolve within the context of a particular environment. One need not restrict one's questions to "how does this strategy do in the tournament?" After all, the details of the tournament change over time. Instead, one can look at a strategy's performance given the other strategies present. This allows a focus on short-term and local successes rather than an overall victory. Thus, predatory strategies that are necessarily short-lived, such as those which wipe out their prey (see Section 3.7), can still be evaluated.

The major disadvantages of genetic algorithms are the time overhead of reproduction and the limitations of the particular genetic code in use. The more possible strategies a genetic code is designed to be able to represent, the more space and time are used. For example, in a 300 round game, a particular human-designed strategy might include, "A defects in round 257 if B defected in round 225". A genetic code that allowed for such fine control of decisions would have to be many thousands of letters long. To implement full control, and allow a decision to be affected by every previous decision, would require 2598 letters just to code the 300th round's decision. Obviously, this is totally impractical. Thus, genetic algorithms must employ some heuristics about what strategies are actually good. Typically, such heuristics include, "a good strategy looks at most at two past moves." This restricts the class of possible strategies significantly, and so one might be concerned that excellent strategies exist in the context of the overall problem but not in the simulation simply because the simulation does not allow the players to examine a sufficient number of moves. However, there is reason to believe that strategies with this limitation are no worse than those without it. For games with random numbers of rounds like the ones used here, almost none of the strategies which did well in Axelrod's tournaments actually looked back farther than two rounds (Axelrod, 1997). Thus, it is unlikely that the set of strategies which would do well looking back an unlimited number of rounds would be much different than those that do well with the limitation.

In Axelrod's first two public tournaments, strategies played against one another for a fixed number of rounds. In both setups, TIT FOR TAT won because of the particular makeup of the strategies involved. However, this is a problematic result. Many of the strategies which contributed to TIT FOR TAT winning did not do particularly well, and would die out in any ecological or evolutionary tournament. Given this, TIT FOR TAT should not actually have been the best strategy. After all, with a fixed number of rounds, one should always defect on the last round. If all players defect on the last round, one should always defect on the second-to-last round. This in turn affects one's options for still earlier rounds (Sandholm, 1996). That this result was not observed points to a significant problem in tournaments with only a few given strategies.

The issue of whether a tournament should be fixed-length or variable-length depends on what one is trying to model. After all, many human interactions are for a set number of episodes. For example, contractors may be collaborating on one multi-step project, and may be unlikely to collaborate again. Other interactions are open-ended in number. However, the ability of strategies to evolve is nearly always part of one's model.

People can almost always change their tactics to respond when they do poorly. Organisms do evolve, and so on. It is rare that stable strategies actually model cooperation more effectively than evolving strategies. In particular, modeling coevolution can only be done when at least some changes in strategies are possible, because evolution requires constant change (Darwin, 1859, Chap. 2).

3.6 Predators/prey, "niceness", exploitation/predation

The customary difference between so-called "nice" and "non-nice" strategies is that the "nice" strategy will never defect first. This is an inconvenient distinction because of the presence of imperfect information, as well as because of the fact that in some games all strategies defect on the first move. A more useful characterization is given by the term "exploiter" (or "predator"), namely "any strategy which is different from many of its neighbors, which is successful because of its differences, and against which neighbors do more poorly than they do against one another." All three factors are necessary for a useful definition. In a colony (a cluster) of frequent defectors, defection is part of their mode of operation; one cannot distinguish between predators and their prey (the neighbors that are doing more poorly, to the predator's benefit). If a defector is not successful in a particular environment, it will not survive or reproduce; it will therefore have minimal effect on its neighbors' evolution. If a strategy is different than its neighbors but helps them, it cannot be said to be exploiting them; instead, it is simply an extra-cooperative entity that can itself be taken advantage of. Note that a predator can help its neighbors indirectly, because it can discourage more dangerous predators from gaining a foothold.

An additional benefit of spatialization for study of predator-prey interactions is that it allows strategies to tailor themselves more closely to their particular neighbors. If an ALWAYS DEFECT is located in a local environment of ALWAYS COOPERATE, it can do fairly well in that environment until it is faced with strategies of different types, or until those cooperators change their respective strategies somewhat. This sort of behavior also occurs in a more global venue, but it is more pronounced when the play of strategies is localized. When a small island (a cluster) of strategies is created, the inhabitants are likely to have common idiosyncrasies (patterns of play), because many of the inhabitants of the island probably have common ancestors. These idiosyncrasies can be exploited by other strategies that behave as "predators", and this causes the inhabitants (strategies) to alter their behavior (presumably by reproduction with mutation and evolutionary pressure) to avoid loss of fitness. The alterations in behavior evolved will not be a general response to exploitation, necessarily; it will be directed only at the particular predator nearby, and may or may not serve to protect the strategy from various other potential predators.

Sometimes strategies such as TIT FOR TWO TATS do better than TIT FOR TAT in a tournament. TIT FOR TWO TATS will defect only after the opponent has defected twice in a row, and may do very well in certain tournament environments. It can also be exploited by a predatory strategy that defects every other round. This means it lacks stability (i.e. the ability to survive) in an ordinary genetic tournament. If TIT FOR TWO TATS is being exploited, a strategy like TIT FOR TAT will begin to do better in comparison. Then, once TIT FOR TAT becomes predominant, the exploiters do poorly. It can be shown that no isolated strategy surrounded by TIT FOR TAT can do better than these TIT FOR TAT neighbors (at least in the absence of noise) (Axelrod, 1984, p40). Thus TIT FOR TAT is a somewhat stable strategy. It cannot be displaced once it becomes the predominant strategy in a local area, unless multiple new strategies enter simultaneously.

3.7 Imperfect information

Imperfect information exists when strategies are given incomplete or incorrect reports about past moves. In the situations examined in this paper, imperfect information (or noise) implies that a strategy is occasionally told its opponent's moves incorrectly. A noise level of $p\%$ in a round means there is a $p\%$ chance that any move in that round will be reported incorrectly to the opponent; neither strategy is informed that this is the case. Other (Axelrod, 1997) implementations of noise are discussed in the next section.

When strategies are told their opponent's moves incorrectly, they may take it that their opponent has defected when in fact, it has not. If this happens to a TIT FOR TAT strategy facing another TIT FOR TAT, the first will retaliate, triggering a retaliation in turn, and resulting in a series of alternating sucker's payoffs and betrayer's payoffs. This series is not as desirable as cooperation, and so more forgiving strategies do relatively better. Forgiving strategies such as TIT FOR TWO TATS (which defect after two simultaneous defections by its opponent) suddenly do have an advantage over TIT FOR TAT, because they can avoid costly mutual defections. Of course, as forgiving strategies become more common, it may become profitable to sometimes defect unprovoked. Thus, there may be more opportunity to exploit one's neighbors in the presence of imperfect information.

Use of imperfect information makes exploitation more common. When TIT FOR TWO TATS is a prevalent strategy, players that occasionally defect and then return to cooperation can benefit. They are not necessarily punished¹ for their defection, and thus manage to obtain a large payoff from a unilateral defection without reducing future payoffs. Such exploitative behavior is much more likely to be profitable (and thus more commonly found) when noise levels are high.

4. USES OF THE SIMULATION, MITOCHONDRIAL GENES, SPECIATION

The Complexity Cooperation (Axelrod, 1997) describes tournaments that explore refinements to his theories regarding the Prisoner's Dilemma. There, Axelrod examines the idea of imperfect information principally to explore the issue of contrition. In many situations, we feel bad when we wrong another accidentally, and we apologize. This frequently translates into permitting one's counterpart in some enterprise (whether business or war) to retaliate (i.e. to defect) without then retaliating in turn. This is what is meant by contrition, for the purposes of Prisoner's Dilemma simulations.

Thus, Axelrod allows a strategy A to know when its actions are misreported, so that it can act differently when the other player B responds to A's actions (intentional or otherwise) than when B defects for no apparent reason. Additionally, Axelrod uses an ecological tournament (no mutations) rather than an evolutionary one, which offers him more control over the strategies in play. He finds that contrition is a trait that tends to increase net payoff and that the ability to apologize (i.e. to exhibit contrition) for what might appear to be a betrayal is important whether or not the betrayal had been intentional. One might well draw parallels between the recent disagreement between China and the US over an aircraft collision and the value of apologizing regardless of fault.

On the other hand, political theorists are not all agreed on the value of contrition for nations. This is because contrition can be exploited, as nations find pretexts for any hostile actions. Thus, for the US to offer compensation to China for its fighter might defuse tensions in the short run, but might make China more likely to create tensions in the future.

¹ The idea of a strategy suffering one or more defections in response to a defection on its part can be regarded as "punishment".

If, as appears to be the case in certain types of conflicts at least, people almost invariably create a pretext when choosing to take advantage of others, then it can become impossible to tell when conflicts are due to mistakes. Additionally, many organisms have not developed the ability to apologize. Thus, for certain types of interactions, it makes more sense not to tell players when noise has caused their action to be misreported. This is the case in the model used here for predator-prey coevolution, since contrition is frequently not part of such interactions. Instead, one is unable to distinguish between a mistake (an apparent defection caused by noise), a response to one's own mistake, and a betrayal (an unprovoked defection). This lack of information makes exploitative strategies more likely to do well than in Axelrod's tournaments.

In most applications, when genetic algorithms are used, it may not be immediately apparent how to evaluate an algorithm's fitness. However, for playing the Prisoner's Dilemma, the algorithm's score (fitness) can be used nearly unmodified to evaluate performance. However, a score might be considered "good" at one time but not at another. If most strategies habitually defect, the average score will be much lower than if most strategies habitually cooperate. Thus, if the death rate is calibrated such that in cooperative populations a few members die in every round, this fact will cause less-cooperative populations to quickly become extinct. The solution used here is to cause players to lose some small amount from their fitness score, in proportion to the number of games they are likely to play. Meanwhile, because of the algorithm used for reproduction, the likelihood of an empty slot being filled increases (at a slower rate of increase than the rate of decrease in fitness) with the number of players near that slot. The result of these two factors is that as population density (the proportion of the slots on a section of board that are filled by living players) increases in an area, the number of strategies that will die will increase more quickly than the number of strategies that will be born. Therefore, a region of cooperative strategies will develop an equilibrium¹ number of individuals reflecting a higher population density than a region of uncooperative strategies; both regions will, however, find some equilibrium where new strategies are continually being born and dying at the same rate.

4.1 Reproduction

Both sexual and asexual reproduction was examined. Asexual reproduction is the simplest way to implement a genetic algorithm: a strategy that does well is permitted to reproduce more often, and so has more children. These children are identical to their parent, unless mutations occur. In sexual reproduction, two parents both contribute to the genetic composition of offspring. This means that the child will have characteristics of both, which has three major effects. First, there is more diversity, since moderately good strategies are still able to reproduce nearly as well as the best strategies. Second, two advantageous characteristics developed independently by different members can combine, and be passed to offspring. (Note that while the act of sex is expensive in terms of energy cost which reduces the rate of production of offspring (Tyler, 2000, Chap 4), we have not yet taken this key element into our simulation.) Thus, while there are reasons to suppose that sexual reproduction would lead to faster evolution than asexual reproduction and reasons to suppose the reverse, it is almost guaranteed that sexual reproduction will lead to more diverse individuals in a population.

¹ Technically, true equilibrium only exists if the cooperativeness of strategies is unchanging over time, but for each round of a game with evolution, a particular equilibrium point exists, so birth and death rates behave as if an equilibrium existed.

4.2 Mitochondrial genes and speciation

In a model of predator-prey coevolution, one aspect of sexual reproduction may be extremely undesirable: that a predator and prey might reproduce sexually. There will not be stable populations of the two strategy classes coexisting if those strategy classes go about interchanging their genetic material. To prevent this, "mitochondrial DNA" has been introduced. For each gene (i.e. element) in the initial genetic sequence used [see next section], another "marker" gene was added. This marker gene determines whether its corresponding initial gene is "mitochondrial", that is, required to appear in the child. If so, the relevant pair of genes automatically show up in one's offspring. Thus, a mitochondrial gene which contributes to a pattern of exploitation in play would automatically produce mitochondrially-exploiting children regardless of its partner. If both parents have mitochondrial genes whose values (see Section 5) are separated by a significant quantity, the two strategies are not permitted to reproduce. They have thus become two separate species. However, in order to allow an isolated mutant to propagate itself, we allow strategies to have some chance of reproducing asexually. This is useful if a player is representative of an entire population of organisms (whose individuals sometimes breed within their population and sometimes do not), or represents an organization that may share information if others are present to receive it, but that can also reproduce on its own.

Once a gene has become mitochondrial (its corresponding marker gene becomes set to "True"), it has improved chances of spreading (mitigated, of course by the possibility that the gene negatively affects the organism's fitness). After all, it has double the normal chance of being passed on. Of course, once differing mitochondrial genes have been developed (via mutation), speciation occurs. One begins to have organisms that cannot breed with one another. Organisms without mitochondrial genes are more likely to breed since they can breed with anyone. However, the offspring need not share their lack of mitochondrial DNA, and so this advantage is not necessarily inherited. Once speciation occurs, it becomes possible to have groups of organisms coexisting but not reproducing together, and thus predator-prey interactions become more likely to occur.

However, it was found that even with mitochondrial DNA, predators (non-nice exploitative strategies) do much more poorly in an environment involving sexual reproduction. This is because more diversity makes it easier to respond to predators. After all, the predator's mode of exploitation must be tailored closely to a weakness in the prey, and diverse prey makes a useful mode of exploitation less likely to develop.

Additionally, a large diverse genetic pool can evolve effective strategies faster than a smaller group. If all individuals have even chances of mutating, the chance of someone mutating a beneficial adjustment and then out reproducing its neighbors, producing a population of better-evolved strategies, is higher if there are more individuals. This aspect is true of both sexual and asexual reproduction. However, in sexual reproduction, creation of a useful new strategy can take place very quickly. Since new players' strategies are combinations of their parents' strategies, a combination of two useful mutations can occur if both arise spontaneously separately. In asexual reproduction, in contrast, both such mutations would have to spontaneously arise in the same lineage of players. Therefore, in sexual reproduction a large group of strategies has a significant advantage over a smaller one, because there is likely to be more genetic diversity. The diversity is advantageous in either sexual or asexual reproduction, but because aspects of different strategies are combined in their children, the diversity is particularly useful if reproduction is sexual. Usually predators are the smaller group, since if prey is being exploited they won't do well unless they can balance this exploitation with considerable successful cooperation with

other prey. So sexual reproduction usually helps prey evolve faster than predators.

Finally, in a situation of sexual reproduction, favorable characteristics are easily shared. If a mutation is developed in an asexual population, it spreads at a rate determined by the death rate of the individuals that are mutating (to make room in the grid for progeny), and the comparative birth rates of the mutant and non-mutant individuals. If the mutation is developed in a sexual population, the increase in prevalence of the mutation in each generation is up to two times the rate in an asexual population. If two mutants breed with one another, this is as effective as in an asexual population at increasing the proportion of mutants. If a mutant breeds with a non-mutant, the non-mutant and mutant genes may be equally likely to show up, in which case this is still only as effective as asexual reproduction. However, if the non-mutant children do poorly, or if the mutant gene is mitochondrial, the mutant gene will spread up to twice as fast (since only one of the two parents really needs to have it). Now, double the rate may not sound like much, but this is double the rate every turn. So if a mutant gene spreads at 5% a turn for 30 turns, it will be present in 4.3 times as many individuals as it was to begin with. Had it spread at 10% a turn for those 30 turns, it would be present in 17.4 times as many. Since growth is exponential, anything promoting growth has greatly magnified effects over time.

The one mitigating factor in predators' favor given sexual reproduction of prey is that the individuals with the greatest pressure to spread an anti-predator mutation are the individuals closest to the predator. However, these individuals probably still do worse than prey without the mutation who are protected from the predator only by distance, because distance is frequently the best defense. This means that there can be situations where the presence of predators results in the death of nearby prey, but these slots are filled more often by distant easy prey than by the tougher nearby mutants.

Note that some believe sexual reproduction was developed in nature to help against predators or parasites (Linde et. al, 2000).

5. THE GENETIC CODE

The genetic sequence of a strategy is specified as follows. There are 43 genes, each of which contains a floating-point number x . The first 21 genes also contain a tag. For genes $m=[0] - [20]$, this number $x[m]$ represents the chance of cooperating; the tag determines when the gene is active (is used). One of these 21 genes is active on each round. The active gene is the gene whose tag corresponds to the previous plays made. The tag consists of four letters P1, P2, P3, and P4, each of which may take the values $\{\emptyset, C, \text{ or } D\}$, representing the empty value (no previous play), a play of "cooperate", and a play of "defect".

P1: denotes the strategy's move 2 turns ago.

P2: denotes the opponent's move 2 turns ago.

P3: denotes the strategy's move 1 turn ago.

P4: denotes the opponent's move 1 turn ago.

For example, a gene with the tag " $\emptyset\emptyset CD$ " is active if there was no turn 2 turns ago (i.e. it is turn 2 of a game), the strategy had cooperated on its first turn, and the opponent had defected on its first turn. The active gene's x -value determines the play. If $0 \leq x \leq 1$, x is the probability of cooperation on that play, and $1-x$ is the probability of defection. If $x > 1$, the play must be "cooperate"; if $x < 0$, the play must be "defect". In Table 1, the tags of genes [0] - [21] are shown. From these tags we see that gene [0] determines the first move, that one of the genes [1] - [4] will be active on turn 2 and will determine the second move, and that on each subsequent round, a gene [5] - [20] will determine the play for that round.

Table 1: Tags for genes [0] - [20]

[0]	øøøø
[1]	øøCC
[2]	øøCD
[3]	øøDC
[4]	øøDD
[5]	CCCC
[6]	CCCD
[7]	CCDC
[8]	CCDD
[9]	CDCC
[10]	CDCD
[11]	CDDC
[12]	CDDD
[13]	DCCC
[14]	DCCD
[15]	DCDC
[16]	DCDD
[17]	DDCC
[18]	DDCD
[19]	DDDC
[20]	DDDD

Gene [21] represents the mutability of the sequence. Its x -value is a floating-point number between 0 and 1. The greater the value of $x[21]$, the more likely children of the sequence are to differ from their parent, and the more significant any such mutations will be. Each gene $[m]$, where $22 \leq m \leq 42$ is a marker gene for gene $[m-22]$, and specifies whether that marked gene is mitochondrial. For example, $x[23]$ determines whether $x[1]$ is mitochondrial. A value $x[m] > 0.5$ denotes that $[m - 22]$ is mitochondrial, while a value $x[m] \leq 0.5$ denotes that it is not.

5.1 Sample sequences

A genetic sequence of plays in the absence of noise is illustrated in Table 2. The first column is the gene number, the second is the associated value of x . The third column contains tags, and the fourth contains comments. This model example was composed so as to show all the interesting frequent occurrences. The amount of historical information (described in the next paragraph) present in this example is much larger than in any typical strategy.

Sometimes genes exist which are never used. For example, since $x[0] \leq 0$, the strategy never cooperates on turn 1. Thus the value for $x[1]$ is irrelevant, since [1] becomes active only if the strategy cooperated on the first turn. Genes which are never read, but which have relevant data, are called legacy genes. They often indicate that ancestors evolved that gene in response to environmental pressures, but that the current strategy happens not to be in need of it.

Table 2: Sample genetic sequence S_1 without noise

[0]	-0.1	$\emptyset\emptyset\emptyset\emptyset$	Strategy begins by defecting on round 1.
[1]	2.2	$\emptyset\emptyset CC$	See note A.
[2]	-0.1	$\emptyset\emptyset CD$	See note A.
[3]	-1.1	$\emptyset\emptyset DC$	See note A.
[4]	-1.8	$\emptyset\emptyset DD$	If both initially defected, strategy defects on round 2.
[5]	1.5	CCCC	Two rounds of mutual cooperation lead to more cooperation.
[6]	-1.7	CCCD	A defection after cooperation leads to defection.
[7]	0.6	CCDC	See note B.
[8]	0.3	CCDD	See note B.
[9]	0.9	CDCC	See note C.
[10]	0.5	CDCD	See note C.
[11]	1.0	CDDC	See note D.
[12]	-0.6	CDDD	If both defect, defects again.
[13]	0.9	DCCC	See note E.
[14]	1.4	DCCD	If opponent takes revenge for S_1 's unilateral defection, allow this.
[15]	-0.0	DCDC	See note F.
[16]	0.6	DCDD	See note F.
[17]	2.0	DDCC	See note G.
[18]	1.5	DDCD	See note G.
[19]	1.1	DDDC	See note G.
[20]	0.9	DDDD	See note G.
[21]	0.3		Some reduction in mutability from conventional starting value, 0.5.

5.2 Notes A-G and comments for the sequence S_1 in Table 2

Note A: Suppose that strategy S_1 is playing others that defect on round 1. Its values for genes [1]-[3] are in some sense unimportant, since they won't become active. After all, neither S_1 nor the strategies it is likely to play are likely to cooperate that first round. Nevertheless, these genes convey important historical information. Some of the strategy's ancestors clearly used to cooperate on the first turn, but these were overtaken by other strategies that defect on the first turn. These cooperative strategies would cooperate only with each other, and they died out because they did poorly against (more numerous or otherwise more fit) initial defectors. Had the cooperators been more numerous, they could have taken over.

Note B: S_1 never defects after mutual cooperation, so there are no evolutionary pressures on these gene values. The values then vary randomly over time. Random variations are unlikely to cause these x-values to change much from their initial values, although this can occur.

Note C: There is some pressure to cooperate after S_1 has been unilaterally defected against, and then cooperated. This is because the only time S_1 will cooperate after a unilateral defection against itself is after it had unilaterally defected against its opponent unprovoked. However, this can also be taken advantage of, as the opponent gets two unilateral defections for its one. Thus, there should be less incentive to cooperate if the opponent had defected that round.

Note D: If S_1 and its opponent continue alternating stabs (i.e. unilateral defections), neither will do particularly well. A strategy may obtain higher fitness values if it can break out of such a cycle by initiating cooperation.

Note E: If S_1 defects while the opponent cooperates, S_1 (playing mostly TIT FOR

TAT) cooperates. Since both cooperated, a cycle of mutual cooperation has probably begun; this gene allows the cycle to continue.

Note F: S_1 should not defect after a mutual cooperation, so there should be no evolutionary pressure to affect these scores. It is of course possible that ancestors' behavior has shaped these, or that they randomly moved to their current values.

Note G: It might seem initially strange that after a mutual defection, regardless of the next turn, the strategy will always cooperate. However, mutual defections can lead to very costly consequences. There are predators that can take advantage of this sort of unresponsiveness, but until these predators force S_1 to change this term, the ability to reach cooperation is quite valuable. Additionally, recall that the first turn is always mutual defection; this policy allows the third move to be cooperation even if the strategy encounters other strategies with different second-turn moves. Given this, it creates some pressure for the second turn to be defection, since any choice will be forgiven (to forgive a defection is to return to cooperation).

In our simulation, each gene's x -value is relevant to 2 places after the decimal point, so a 0.6 might be anywhere from 0.55 to 0.64. For purposes of understanding patterns in the data the additional decimal place is only sometimes helpful, and is not shown in Tables 2, 3, or 4.

Sometimes x -values below 0 or above 1 will be found. This can be beneficial, because mutations happen. If a strategy is best served by always defecting in a particular circumstance, and has a gene value of 0, some of its children may have positive gene values, and so they will sometimes cooperate. Therefore, strategies with negative gene values will have more successful children. So there is some evolutionary pressure to have very low or very high gene values, even though the individual with those genes behaves (i.e. plays) no differently than one with a strategy with x -values closer to 0 or 1 respectively.

To some small extent, the extremeness of a gene's x -value is indicative of its importance. If a gene has only a small effect on fitness, those children with a somewhat less extreme gene may do nearly as well as their parents, so there is not much pressure to increase the x -value above 1 or reduce it below 0. If it has a large impact, the evolutionary pressure is much greater, and the final value can often be more extreme.

The mitochondrial DNA is not shown in Table 2, because there is no real pattern regarding which gene becomes mitochondrial in non-predators. It is interesting that mutability (gene[21], whose x -value is permitted to vary in the same way as other genes) tends to decrease in value as time passes, but not significantly if reproduction is sexual. When reproduction is asexual, mutability frequently becomes much smaller. This is probably because with sexual reproduction the children are not often similar to their parents anyway, and thus mutations are not as frequently disadvantageous. Of course, in either case this is balanced at some point by the advantages of mutability, such as increased resistance to predators and increased ability to take advantage of novel strategies.

In the presence of noise, conditions that are appropriate for non-noisy strategies frequently result in the extinction of all strategies. This is because significantly more defection occurs when some co-operations are seen as defections, and so the average scores are much lower. It is therefore necessary to make the death rate much lower to make up for this. Doing so allows individuals to survive even if they do not quickly evolve particularly cooperative strategies. In fact, in any sufficiently long tournament global extinction is the expected result. After all, if all the players have an independent, small, finite chance of dying, there is a non-zero chance that all strategies will simultaneously die. This chance is minute for any given round, if fitness levels are sufficiently high.

With noise, two major types of strategies tend to become widespread. In certain games, especially with very low noise levels, slightly forgiving variants of TIT FOR TAT seem to emerge. Such strategies are much like those versions of TIT FOR TAT which arise in a noiseless environment, but have some slight chance of cooperating in response to defection, and thus increased ability to avoid runs of mutual defection. A typical example, where each move has a 2% chance of being misrepresented to the opponent, is shown in Table 3. This example, like that in Table 2, is a model example. For all genes whose average x-values do not differ significantly from those found in the previous simulation, the values shown are the same as those shown in Table 2. This allows easier comparison of the two tables. With worse information (noise levels above 4%), as well as in many games with small noise levels (~ 1-2%), a somewhat less cooperative strategy becomes dominant (i.e. most prevalent). In such cases, unless the death rate is set very low, extinction of all strategies frequently occurs. However, when the death rate is set low, it is frequently possible for strategies which alternate defection to become prevalent.

Table 3: Sample genetic sequence, S₂ with noise of 2%

[0]	-0.1	
[1]	1.0	See note H.
[2]	-0.1	
[3]	-1.1	
[4]	-1.8	
[5]	1.0	See note I.
[6]	0.0	See note I.
[7]	1.5	See note J.
[8]	1.2	See note J.
[9]	1.1	See note I.
[10]	-0.3	See note I.
[11]	1.0	
[12]	-0.6	
[13]	0.9	
[14]	1.4	
[15]	-0.0	See note K.
[16]	-0.1	See note L.
[17]	2.0	
[18]	0.8	See note M.
[19]	1.2	See note M.
[20]	0.9	See note K.
[21]	0.3	

5.3 Notes H-M and comments for the sequence S₂ in Table 3

Note H: Strategies which initially cooperate do even more poorly with imperfect information than with perfect information, both because initial defections sometimes occur, and because strategies are more likely to have defections later. An initial defect response is not overly costly. A legacy gene relating to initial cooperation is therefore less likely to be found in strategies, and if one is found it's x-value likely to be slightly lower, since there were probably somewhat weaker environmental pressures selecting for it. A gene whose x-value is high is probably strongly selected for: that is, there are (or were) strong environmental

pressures moving its value to one extreme or the other.

Note I: Strategies are likely to cooperate when cooperated with, to defect when defected against, and to forgive on the next round if the opponent cooperated. However, these x-values are not particularly extreme. This means that some of the individuals will occasionally defect unprovoked, that some will not always forgive, and that some will sometimes forgive a defection. This does not strongly affect behavior for the majority of individuals in the majority of rounds (since the values tend to hover close to 0 and 1), but for some children there is a real difference. It is interesting to note that the response to two unilateral defections in a row is always to defect; this is a strongly selected-for response.

Note J: After perpetrating an unprovoked defection, the strategy S_2 will always cooperate. This is probably because it can sometimes get away with an isolated defection, but never with two defections in a row. This is a more extreme value than the possibility of defecting in the first place, which seems strange. After all, what to do after defection can only be selected for after defection is chosen. However, the decision to defect or not receives responses in both directions (bad idea to defect, or good idea), whereas the decision to cooperate after a defection is almost always the right choice.

Note K: Genes [15] and [20] are complementary. If one strategy defects and one cooperates, and then both defect, the next move tends to involve one defecting and the other cooperating (since this brings the next move to the decision coded by genes [18] and [19], usually cooperation). If genes [15] and [20] both cause the strategy to cooperate, the strategy could benefit by defecting instead, since no retaliation would result. However, if either gene typically results in defection, the competing strategy is better off cooperating, since at the cost of this one stab, a string of mutual defections can be avoided in favor of a string of cooperation. Thus, in some games gene [15] tends to be high and gene [20] low; in other games it is the other way around. It would seem "fairer" to a human observer for the initial defector to have to cooperate, but in practice either can evolve. "Fairness" does not necessarily evolve.

Note L: If both defect after S_2 had unilaterally defected, both defect (since x_{12} is negative). However, this is acceptable since after the second set of mutual defections, both will cooperate. It therefore makes sense to do as well as possible after the first defection, since doing so does not trigger more retaliation.

Note M: If one strategy cooperates after both defect, both will cooperate on the following round. This allows the two strategies to continue cooperating, since the x-values of genes [9] and [13] (the two next choices) are both high. This allows a run of cooperation to be reached.

In noisy situations, there tends to be much more variation in eventual strategies. There are a number of reasons for this. First, how well a strategy performs becomes slightly less important since none do extremely well and there are more good strategies than excellent strategies. Second, because it is harder to distinguish from random mistakes, exploitation is easier, and thus there are more predators created and more prey strategies evolved so as to avoid the predators. Third, strategies do not have to be exactly in sync with one another as long as they do not get trapped into long periods of mutual defection. Indeed there is a lower standard of success, because the other players also do more poorly than in a simulation without noise, and so, a lower fitness score is still competitive.

Repeating the simulation shown in Table 3 with higher levels of noise produces results similar to those seen in Table 4.

Table 4: Sample genetic sequence S_3 with noise of 5%

[0]	-0.6	S ₃ always defects the first round.
[1]	-1.9	See note N.
[2]	-0.9	See note N.
[3]	-0.3	See note N.
[4]	0.2	See note N.
[5]	3.4	If cooperation is achieved, maintains it.
[6]	-1.0	If opponent defected, S ₃ will defect next.
[7]	-0.9	If S ₃ defected, it will defect again.
[8]	2.3	If both defect, will cooperate.
[9]	0.2	Shouldn't come up much.
[10]	0.1	Shouldn't come up much.
[11]	-0.6	After defecting, defects again.
[12]	-0.3	See note O.
[13]	-1.0	See note P.
[14]	2.7	Allows alternation of defection-cooperation.
[15]	-0.2	If S ₃ defected, will defect again.
[16]	-0.8	See note O.
[17]	-0.3	See note Q.
[18]	-1.1	Is S ₃ 's turn to defect.
[19]	1.0	Is S ₃ 's turn to cooperate.
[20]	0.4	Doesn't get caught in an endless string of defections.

5.4 Notes N-Q for the Sequence S₃ in Table 3

Note N: S₃ always defects on the first round. The second round is also usually a defection. Since the strategies typically survive by alternating defections, it is sometimes better to cooperate if the other player defects, because then an alternation is begun without first mutually defecting for several rounds. Of course, it is better to be the defector there, so the chances of cooperation are much less than 50%.

Note O: After one cooperates and one defects, then both defect, both will again defect. This then leads to gene [20], which eventually may cause another run of alternating cooperation.

Note P: This shouldn't usually come up, however, S₃ can defect every other round. This means that any vulnerable strategies (those that do not respond directly to defection) will be taken advantage of, will have lesser fitness, and will not breed very successfully.

Note Q: Could just as easily be positive. Mutual defection then mutual cooperation is just as good as alternating stabs. However, in this case, S₃ breaks out of this cycle by defecting. This is because if only one defected, it would have a slight advantage, so the defectors win out.

5.5 Common tactics

It is interesting to note that the most common strategy evolved in the presence of perfect information is "DEFECT, COOPERATE, TIT FOR TAT"¹. This seems strange, since strategies would all be better off if they cooperated initially rather than defecting initially. However, the genetic code must characterize the first two rounds differently from those that follow, because less information is available (the code remembers two rounds in the past). Given this, it is possible to treat the first two rounds differently from later ones. Then, a

¹ That is, defect on round 1, cooperate on round 2, then play TIT FOR TAT thereafter.

large cluster of "DEFECT, COOPERATE, TIT FOR TAT" is not easily invaded¹, since ordinary TIT FOR TATs will find that they do better against one another than the dominant strategy (DEFECT, COOPERATE, TIT FOR TAT) does, but do slightly worse against the dominant strategy than the dominant strategy does against itself. The number of TIT FOR TAT strategies present in a local region would have to be very large for this tradeoff to prove advantageous. Indeed, supposing a payoff matrix such that 4 points were gained against itself, and 1 is lost against the dominant DEFECT, then 25% of a local population would have to be of the invading type for this tradeoff to prove advantageous. The situation is even worse, since reproduction is proportional to the number of individuals present. For TIT FOR TAT to be the fastest-reproducing strategy under this payoff matrix, nearly half the individuals in a local population would have to be TIT FOR TAT. Since this rarely happens, if "DEFECT, COOPERATE, TIT FOR TAT" takes over, it frequently remains dominant.

Of course, if TIT FOR TAT becomes dominant, it is not easily invaded by "DEFECT, COOPERATE, TIT FOR TAT". After all, these invaders will find that TIT FOR TAT retaliates (defects) for their initial defection, and forgives them after their second cooperation; they do only slightly worse than TIT FOR TAT. There will be no more impetus for them to take over than there was for TIT FOR TAT. The reason "DEFECT, COOPERATE, TIT FOR TAT" usually becomes dominant is that it is a better strategy early in the game. When life evolves, defection is usually the best strategy, since many organisms have not learned to respond properly. Many will cooperate in response to defection or defect in response to cooperation; until these die out, defection is rewarded. Thus, the initial defect has an early advantage. Then a TIT FOR TAT commencing at round 3 response is evolved, because this eventually becomes a worthwhile strategy (once the unresponsive strategies are removed), but there is little environmental pressure to change the initial defection to cooperation.

Then in order to create an environment where TIT FOR TAT is the dominant strategy, it is necessary to alter the starting conditions. If strategies are initially generated completely at random, then early exploitation is developed, which gives way to responsive² and potentially cooperative strategies. There is no reason, however, to suppose that initial conditions ought to be totally random in this way. Rather, the initial conditions should reflect those of the scenario being explored. If organisms or organizations are assumed to begin in a state of cooperation, it is not difficult to create strategies that are mostly cooperative, with some variation, and see how evolution proceeds. The state of the board at the beginning of the tournament has a significant role in determining the eventual strategies' first moves during each game, but a much more minor role in determining their later moves in each game.

A strategy's first turn, one might think, should be similar to its behavior on later turns. If a pattern of alternating cooperation/defection arises, this could become a stable system. In our simulation, alternating cooperation and defection is not much worse than full cooperation. Indeed, TIT FOR TAT can easily evolve into a strategy designed to ensure alternating cooperation if both a small chance of defecting after mutual cooperation and a chance of cooperating after mutual defection, are introduced. This outcome is in fact

¹ A cluster of strategies in a region can be "invaded" by a strategy foreign to (i.e. not present in) that region if the new strategy does sufficiently well playing the strategies in the cluster that it survives and reproduces.

² A responsive strategy's moves are consistently dependent on its opponent's previous moves. TIT FOR TAT is such an example, and DEFECT is an example of the opposite.

sometimes seen, both in games with sexual reproduction and asexual reproduction. When imperfect information is introduced, this outcome is seen significantly more frequently. This is probably because imperfect information makes defection much more likely, so strategies that attempt to mutually cooperate do not succeed all the time. Thus the difference between attempting to cooperate and attempting to alternate defection is not very large. If this is the reason, then we would expect changing the payoffs to change the rates of occurrence of cooperation and defection.

Many strategies can be regarded as exploiters. The most obvious would be ALWAYS DEFECT, but exploitation need not be so overt. In Axelrod's [1984, p. 37] initial set of experiments, a clever strategy which did not perform particularly well was JOSS, a strategy which acted like TIT FOR TAT most of the time, but occasionally would defect regardless of its opponent's move. Such a strategy would be likely to perform much better in an environment with imperfect information, because it would face more forgiving strategies.

5.6 Workings of a gene

Gene [0] has the most direct effect on the game, since it is always used and determines the first move of the game. Its contribution to fitness depends primarily on genes [1] - [4]. Suppose there is a commonly occurring strategy in a region of the grid. If that dominant strategy cooperates on the second turn if and only if there is cooperation on the first turn, then it becomes more profitable in the short run to cooperate on the first move, as long as genes [5] - [8] are at least as profitable as genes [13] - [16]. (Which of these two sets is used on turn 3 is determined by play on turn 1). However, while genes [1] - [4] evolve concurrently with gene [0], changes in gene [0] have a larger average impact on fitness than any of the genes [1] - [4]. Whereas a defection on turn 1 (gene [0]) immediately helps the strategy, retaliation on the second turn is only helpful if it does not trigger more retaliation from the opponent. This is why it is more likely for DEFECT to evolve on turn 1 than on turn 2.

6. TOURNAMENT SIMULATIONS

10 tournaments with no noise, low noise, and some noise were conducted on 11x11 grids, using sexual reproduction primarily, and mitochondrial genes. In Tables 5-10, we show the evolution of gene [0] in various contexts. The second column shows the percent of strategies that cooperate after the specified number of turns (first column); the third shows the percent that defect after the specified number of turns; the fourth shows the percent with an x-value in between cooperation and defection. All x-values are initially generated at random uniformly between 0 and 1.

Table 5: Gene [0] with no noise, with sexual reproduction

turn	cooperates ($x[0] > .8$)	defects ($x[0] \leq .2$)	indeterminate ($.2 < x[0] \leq .8$)
0	20.4%	19.3%	60.3%
30	13.2%	45.2%	41.6%
100	2.8%	87.6%	9.6%
500	0%	100%	0%

Table 6: Gene [0] with 2% noise, with sexual reproduction

turn	cooperates ($x[0]>.8$)	defects ($x[0]\leq.2$)	indeterminate ($.2<x[0]\leq.8$)
0	19.7%	19.8%	60.4%
30	10.2%	44.7%	45.1%
100	0.2%	94.6%	5.2%
500	0%	100%	0%

Table 7: Gene [0] with 5% noise, with sexual reproduction

turn	cooperates ($x[0]>.8$)	defects ($x[0]\leq.2$)	indeterminate ($.2<x[0]\leq.8$)
0	20.1%	20.4%	59.5%
30	10.4%	45.3%	44.3%
100	0%	97.5%	2.5%
500	0%	100%	0%

The data becomes much different with asexual reproduction: cooperation on the first turn becomes more prevalent (in the no noise and low noise cases). Indeterminate strategies also become more prevalent.

Table 8: Gene [0] with no noise, with asexual reproduction

turn	cooperates ($x[0]>.8$)	defects ($x[0]\leq.2$)	indeterminate ($.2<x[0]\leq.8$)
0	24.4%	18.4%	52.8%
30	12.4%	23.6%	54.0%
100	12.9%	25.3%	61.8%
500	13.2%	27.7%	59.1%

Table 9: Gene [0] with 2% noise, with asexual reproduction

turn	cooperates ($x[0]>.8$)	defects ($x[0]\leq.2$)	indeterminate ($.2<x[0]\leq.8$)
0	26.2%	18.4%	55.4%
30	31.5%	16.3%	52.2%
100	30.4%	19.5%	50.1%
500	30.4%	21.4%	48.2%

Table 10: Gene [0] with 5% noise, with asexual reproduction

turn	cooperates ($x[0]>.8$)	defects ($x[0]\leq.2$)	indeterminate ($.2<x[0]\leq.8$)
0	19.3%	24.3%	56.4%
30	2.6%	32.8%	64.6%
100	0%	12.1%	87.9%
500	0%	11.4%	88.6%

In the sexually reproducing instances (Tables 5, 6, 7), cooperation on the first turn was eventually eliminated. In other test cases, cooperation was found to sometimes win out over defection; however, no such instances occurred during the 10 tournaments examined. It did appear, however, that in a few of these scenarios, cooperation won in at least one local area for some time. However as the local pockets of initial cooperators and initial defectors met,

the cooperators always lost. Increased noise did tend to make cooperation still less viable.

In the asexually reproducing case, initially cooperative strategies won far more often (the case of high noise excepted). This may be because strategies are more likely to be found near their own children, and so when a gene for cooperation benefits organisms near the organism with that gene, those organisms tend to have the gene in question.¹ It might be interesting to examine genetic sequences encoding for the distance at which the children are likely to be born, and thus the frequency of encounters between children of the same parent.

Data of this sort is significantly less useful for genes other than gene [0]. The prevalence of a particular value for $x[0]$ can be evaluated on its own under narrow assumptions about its relevance: it is used exactly once every game. For a gene such as gene [1], determining in what situations the gene is actually selected for becomes more complex. If $x[0]$ is small, gene [1] never comes into play; gene [1] can only be meaningfully examined if both the strategy in question and some of its neighbors sometimes cooperate. Its value must be examined in strategies whose value for $x[0]$ is sufficiently high, in reference to $x[5]$ vs. $x[7]$ of the strategy in question, of $x[5]$ vs. $x[6]$ of the neighbors with $x[0]$ values which are sufficiently high, and of genes $x[7]$ vs. $x[8]$ of neighbors with $x[0]$ values sufficiently low.

Given this amount of complexity, it becomes essentially impossible to do a generalized appraisal of gene [1] for every type of tournament that can be run. Instead, it becomes important to examine the particular details of a particular simulation, and observe the behavior in that specific environment. Additionally, this allows one to examine specific genetic patterns, rather than solely the proportions found in the aggregate. For example, one can find the size and number of clusters of a particular strategy, rather than just its presence as a percentage of the whole.

As an example of the sort of detail of evaluation that can be performed on a particular simulation, a sample application of Prisoner's Dilemma is studied. This experiment sets up parameters suitable for modeling the behavior of lichen, and examines gene [1] with reference to the other relevant genes in the simulation.

7. APPLICATION: THE EVOLUTION OF LICHENS

Lichen consists of colonies composed of both fungi and algae. The fungi are able to attach to rocks, to extract some minerals and nutrients from rock, and to trap moisture. Algae conduct photosynthesis, and provide energy to keep the fungi alive. Together, the colonies are able to flourish in extremely inhospitable conditions where neither organism could possibly survive on its own. The fungi would starve on bare rock, and winds or drought would destroy algae. Accordingly, the lichen can only exist as a symbiotic colony of both fungi and algae.

To model this scenario as an iterated Prisoner's Dilemma tournament, the tournament board slots are each initially populated with players who are either of a fungal species or an algal species. Algae are only permitted to play games with fungi, and vice-versa. However (recalling our discussion of death rate in Section 4), proximity to an organism of either species continues to lower an organism's fitness level (by the factor 0.1). While a game (involving sufficient cooperation) with an organism of the complementary species can increase a player's fitness to help make up for the loss of fitness due to proximity of the two organisms, no intra-species games are permitted. An organism

¹ This may be a parallel to the fact that in biology, genes for altruism which do not help the altruistic individual still propagate in many populations.

therefore does best when surrounded by individuals of the complementary species, and not by individuals of its own species.

Initial populations: In the initial generation of the board, organisms are equally likely to be fungi or algae. Organisms of either species are generated in equivalent manners, so while x-values for the two species may diverge, there is no reason to keep track of the tags "algae" or "fungi" between tournaments. The initial x-values for each gene are generated randomly between .66 and .99 (thus indicating an initial propensity to cooperate in most circumstances), except for genes [0] and [1] (which are under consideration here, and are therefore generated at random uniformly between 0 and 1), and gene [10] (which is set to 0, since one may see that a higher value for x[10] is never beneficial, and so this choice accelerates evolution slightly). The remainder of development occurs by customary evolution in the tournament.

The payoff matrix in Table 11 is designed as follows: algae and fungi can cooperate by sharing nutrients and food fairly, and can defect by taking more than their share of food and reproducing rapidly. A symmetric payoff matrix was used, with the following values (fitness contributions): suckers receive -.05 fitness; betrayers receive +.05 fitness; mutual defectors get -.02 fitness; mutual cooperators get +.02 fitness. After each round of any game, we fix the chance that the game will end at 1%.

Table 11: The payoff matrix for players A and B

		A's payoff	B's payoff
C	C	.02	.02
C	D	-.05	.05
D	C	.05	-.05
D	D	-.02	-.02

In the simulation, reproduction is asexual. No changes in species occur: all descendants of fungi remain fungi, and all descendants of algae remain algae. Genes have a 60% chance of being passed on unchanged, a 24% chance of being passed on with a small mutation (a normal distribution centered around the current value, with a standard deviation of about .05), and a 16% chance of being passed on with a larger mutation (a normal distribution centered around the current value, with a standard deviation of about .15). The exact value of the standard deviation is dependent on the mutability of organisms, gene [21], which is permitted to evolve like the other genes.

The chance of a game being played between any pair of players is large for next nearest neighbors in the grid, with a rapid drop-off with distance. Nearby individuals have a 50% chance of playing one another; we set this to drop by $.65^d$, where d is the distance between the organisms. Likewise, the chance an empty slot is filled by reproduction of a particular neighbor is proportional to that neighbor's fitness and to $.65^d$, where d is the distance between the neighbor and that empty slot. Thus reproduction is usually fairly local.

The playing grid was set to 11x11, primarily because the time required to run the simulation increases with the square of the area and thus with the fourth power of the length of a side. Increasing side length above 11 significantly increases the time to run a simulation, and was not observed to significantly alter the results of most simulations.

To begin, the evolution of the first move (gene [0]) was observed in 10 tournaments. The average number of cooperators ($x[0] \leq 0.2$), defectors ($x[0] > 0.8$), and indeterminates ($0.2 < x[0] \leq 0.8$) at 0 turns, 30 turns, 90 turns, 270 turns, 500 turns are shown in Table 12.

Table 12: Rate of plays of gene [0] over the course of play

turns	cooperators	defectors	indeterminate
0	.21	.22	.57
30	.22	.26	.52
90	.19	.18	.67
270	.30	.14	.56
500	.24	.27	.49

These numbers are averages, and do not necessarily indicate that evolution did not occur. In each game, strategies surviving past 100 rounds or so tended to be either cooperators/indeterminate, or defectors/indeterminate. In no instances was it found that one species evolved to initially cooperate and the other to initially defect. Therefore, it is reasonable to examine gene [1] in strategies that initially cooperate, because they frequently will encounter other strategies that initially cooperate.

In round 2, only strategies where $x[0] > 0.5$ (those which primarily cooperate) were examined. Only games with sufficient initially cooperating strategies (>30 out of 121) were examined. 10 such rounds were obtained, and the value of gene [1] is shown in Table 13.

Table 13: Rate of plays of gene [1] over the course of play

turns	cooperators	defectors	indeterminate
0	.19	.20	.61
30	.14	.28	.58
90	.24	.17	.59
270	.36	.21	.33
500	.43	.26	.31

It is clear that cooperation becomes an improving feature as time progresses. However, since these results are aggregates, they do not show how an individual game progresses. What occurs, generally, is that (on a local level) the strategies move towards either cooperation or defection, with a few intermediate values on either side. Occasionally, the two species will diverge, with algae moving to cooperation and fungi to defection, or vice-versa. By round 500, this process is complete: each species will have chosen a strategy (although some individuals will deviate from this), and the two species' strategies will usually be the same. To demonstrate how this works in the early game, a few rounds from each of two games are exhibited in Tables 14-16.

In this first game, a result, symmetric between the two species, is seen: both algae and fungi begin by primarily cooperating. This tends to build on itself quickly, with both species progressing towards more cooperation. In Table 14, we show the relevant $x[1]$ values on the board for the symmetric game, during three rounds. Recall that strategies with $x[0] < 0.5$ are not of interest, and so the corresponding $x[1]$ values are not shown.

Table 14: Display of x[1] on the tournament board for rounds 34-36

key: bold face: algae.

blank: dead.

J: x[0] < 0.5

round 34

0.6		0.6	0.6	J	0.8		J	J	
0.6		0.6	0.8	J	J	J			J
0.6	0.6		0.8	0.6	J		J	J	J
	0.6	0.6		J		J	J	0.8	J
0.6	0.6	J	J	J	J			J	J
J	J		J	J	J		J	J	J
J	J					J	J	J	J
J	J	0.78	J	J	J		J	J	J
J	J	J	J	0.6	J	J	0.78	J	J
0.78									
	J	J	J	J			J		
	J	J	0.78	J	0.78				

round 35

			0.6	J	J				
		0.6	0.6	J	J	J		J	J
		0.6	0.6	0.6	0.6	J		J	J
0.8	0.6	J	0.6	0.6	J	0.8	J	J	J
J	0.6	J	J	J	0.8	J		J	J
J	J	J	0.6	0.8	J	0.6	J		J
0.78									
	J	0.6		0.6	J	0.6	J	0.78	J
0.78	J	0.6	J	0.6	J	0.8	0.8		J
J	0.6	J	J	0.6	J	J	0.78		J
J	J	J	0.8	J	0.8	J	J		
			J	J	0.6	J	J		

round 36

	0.6		0.6	J	0.6	J			J
0.6	0.6	0.6	0.6	J	J	0.8	0.78	J	
			0.6		0.6	J		J	
0.6	0.8	0.6	0.6	0.6	J	0.6	0.8	J	J
J	0.6		J	0.6	0.8	J	J	J	
J	J	0.8	0.6	J	0.6	0.6	0.8	J	
	0.8	J		0.6	J	0.6	J	J	
J	0.8	J	0.8	0.8		J	J	J	
J	0.8	0.8	J	0.6	J	0.8	J		
J	J	0.8	J	J	0.8	J		J	J
	J	J	0.6	0.8	J		J		

In Table 15, we show three rounds of the relevant x[1] values on the board for an asymmetric game. The asymmetric result at this point: algae cooperate or defect, while fungi

only cooperate. It is interesting that the asymmetry is so marked: in cases where it is less marked, the tendency is to converge on mutual cooperation or mutual defection.

Table 15: Display of x[1] on the tournament board for rounds 34-36

key: bold: algae.

blank: dead.

J: x[0] < 0.5

round 34

0.96	0.1		0.1	0.96	0.1	0.1	0.1		0.1	0.1
0.96	0.96	1.0	1.0	0.96	0.96	J	1.0	0.1	0.96	0.1
1.0	0.1	0.1	0.1	0.1	0.1	0.1		1.0	0.1	0.1
0.1	0.96	1.0	1.0	0.96		0.1	0.1		0.1	1.0
0.96		0.96	0.96	0.96	0.96	1.0	0.1	0.96	1.0	0.96
0.1	0.96		0.96		0.96	0.96	0.96	0.1	0.1	1.0
0.96		J	0.96	0.96	0.1	J		0.96	0.1	0.96
0.1	0.96		0.1	0.1	0.96	0.96	0.96	0.96	0.96	
0.96		0.96		1.0	0.96	1.0		0.96	0.96	0.96
	1.0	0.1	0.96	0.1	0.96	1.0	J	0.1	0.96	
0.1	0.96	0.1	0.1	0.96	0.1	0.1	0.96		0.96	

round 35

0.96	0.96		0.1	0.96	0.1	0.1	1.0			
0.96	0.96	0.96	0.1	0.96	0.1	0.96	0.1	0.1	0.1	0.1
1.0	0.1		J	1.0	0.1	0.96	0.1	0.1	0.1	
0.1	0.96	0.96	0.1	0.1	0.96	1.0		0.96	0.1	0.1
	0.96			J	0.96	0.96	0.96	0.96	0.96	0.1
0.1	0.96	0.96	0.96	0.1		0.1	0.96	1.0	0.1	1.0
0.96		0.96	0.96	0.1	0.1	1.0		0.96	0.1	0.96
0.1	0.1		0.96	0.1	0.96	0.1	0.1	0.1	0.96	1.0
0.96	0.1	0.96		0.1	0.96		0.96	1.0	1.0	0.96
0.96		0.96	1.0	0.1	1.0	0.96	0.96	0.1	0.96	0.96
	0.96	0.1		0.96	0.96	0.1	0.96	0.96		0.96

round 36

	1.0	0.96		J	0.96	0.1				
0.96	0.1	0.1	0.1	0.96	0.1	0.96	0.1	0.1		0.1
0.96	1.0	0.96	0.96	0.1	0.1	0.1	0.1	0.1	0.96	
	0.96		0.1		0.96	0.1		0.96	1.0	
0.96	0.1	0.96	1.0	0.1	0.1	0.96	J		0.96	0.1
0.1	J	0.1	0.96	0.1		0.1	0.96	0.1	0.1	0.96
0.96	0.96	0.1	0.96	0.1		1.0	0.96		0.1	0.96
0.1		0.96		0.96		1.0	0.1	0.1	0.96	J
0.96	0.1		0.96	1.0		0.1	0.96		1.0	
	0.96	J		0.1	J	0.96	J	0.1	0.96	
J		0.1	0.96	0.96	0.1	0.96		0.96		

In order to demonstrate whether divergence of the two species' strategies is rare, 500 games were played using the same parameters as above. For each game, 400 rounds are played; only games in which at least 30 strategies have $x[0]$ values above 0.5 are counted. For these games, it is determined whether the majority of fungi have $x[1]$ values below 0.3 (called fungi defect) or above 0.7 (called fungi cooperate), and whether the majority of algae have $x[1]$ values below 0.3 (algae defect) or above 0.7 (algae cooperate). The compilation of these results is shown in Table 16.

Table 16: Number of games (out of 500) showing each type of result

#games	Result
84	Extinction
143	Less than 30 strategies have $x[0] > 0.5$
41	Fungi defect; Algae defect
45	Fungi defect; Algae cooperate
52	Fungi cooperate; Algae defect
135	Fungi cooperate; Algae cooperate
179	At least one defects
232	At least one cooperates

This indicates that cooperation is more likely than defection, and that there is also a tendency towards symmetry. At first this seems a bit strange, since there are more of each of the two types of pairs in which one species cooperates and one defects than of mutual defections. This seems a result of a tendency towards more cooperation than defection.

However, there are 41 games involving mutual defection, and 135 involving mutual cooperation. If this were a purely random result, it would mean that there is a 15% chance of mutual defection¹, and thus a 39% chance a particular strategy defects². Then, we would expect about 65 each of the only-one-defects situations, and 102 of the both-cooperate. The fact that there are only 66% the expected number of mixed situations, and 133% of the expected number of mutual cooperations suggests that there is also a tendency to do what opposing strategies are doing.

Information sharing does not create this tendency³. Since reproduction is asexual, so there is no genetic communication between algae and fungi. Further, there is no reason encoded in the genes or the Prisoner's Dilemma games themselves for why symmetry is favored. However, one possible explanation is that the fact that algae and fungi need one another tends to favor symmetry. Algae exploiting fungi is favorable in the short run (which is why defection does well initially as a strategy), but this reduces the number of fungi, which harms that algae population in the long run because comparable numbers of both is important for a colony to survive.

In order to further test the hypothesis that lack of excessive competition is primarily due to the undesirability of harming members of the other species, it would be interesting to run simulations in which one population had a different mutation rate than the other. In this

¹ Mutual defection occurs in 41 of the 273 (=500-84-143) evaluated games.

² The chance that both defect is the square of the chance a particular one defects.

³ The transfer of genetic information from one player to another (or to the offspring of another). Sexual reproduction is one form of information sharing, as is the exchange of DNA between bacteria.

way, we could look further at the reasons why exploitation does not always occur in this sort of symbiotic relationship. For example, we might find that when one species has a very low mutation rate and thus little ability to protect itself from exploitation, the other species acquires various predatory behaviors, but does not necessarily do best by being overly exploitative. Rather, it might be the case that when the rapidly evolving species defects too frequently, it annihilates the stable species in its region and harms itself in that way. However, this is only one possibility. It might instead be the case that the optimum level of exploitation in an unbalanced experiment is much higher than is found in our balanced experiments, and that there is therefore a different reason why symmetry tends to evolve. Perhaps a gene that rewards symmetry on the second round is also frequently helpful in later rounds. For example, it is possible that $x[5]$ is usually high, so that mutual cooperation on both the first and second rounds invariably begins a cycle of long-term mutual cooperation. Gene [5] is evolved more quickly than comparable genes whose effects during a particular game are not symmetric (the same for both species). Such a reason for symmetry would have to do with the simplicity of the genetic code, and its effect might or might not be lessened by a code that went back three generations, for example. To test this hypothesis, we would need to explore alternate genetic codes. It might be interesting to inform our model by examining the genetic memory of algae and fungi, and to determine experimentally their respective abilities to respond to historical events. If these abilities are not symmetric, that could be the basis of a somewhat different model for coevolution.

8. CONCLUSION

The parameters defined by this implementation of the Iterated Prisoner's Dilemma form a powerful and versatile tool for modeling a wide variety of situations involving cooperation and defection (exploitation). However, this very versatility makes it important that a model be appropriate (the particular parameter values) for the particular situation of interest. There are a large number of means of exploitation that can arise, and most are useful or likely to arise only under extremely narrow sets of conditions. It is therefore very important to know the details of the situation under consideration as intimately as possible. For example, in examining the evolution of even two genes in algae-fungi interactions in the context of lichen, it is important to consider an enormous number of factors. Thus, while this sort of simulation provides a useful framework, it needs to be supplemented with careful specification of the actual conditions surrounding the interaction in question. The more closely the model resembles the actual situation, the more likely it is to provide useful clues as to what behavior is actually in an entity's best interests. This can be used as a starting point for exploring why behavior might not reflect what self-interest seems to require.

The methods used to model coevolution in lichen are applicable to a wide variety of other situations. Any numbers of biological contexts suggest potential modifications appropriate to a simulation. Running speed among deer and wolves, fever symptoms in humans, and vertical growth in trees all appear to involve localized coevolution with imperfect information. A particularly interesting example arises from the fact that nestlings compete with one another to be fed by crying. To simulate this, one might allow a pair of organisms to have multiple offspring at once, and have all Prisoner's Dilemma games be played between the newborn strategies. If another species were simultaneously introduced with the ability to play Prisoner's Dilemma games with the offspring of the original species, we might expect that this new species would be more likely to defect. After all, it may share few genes with the organisms, against which it generally plays and might be less likely to

demonstrate altruism. Such a setup might be useful for modeling the behavior of cowbirds, which lay their eggs in other birds' nests, and whose chicks tend to be quite loud and successful at demanding food from their host "parents".

A non-biological simulation is the societal development of "police" officers. When noise levels are high but predators are scarce, strategies that are cooperative even after apparent defections may do quite well, since they can avoid generating mutual defections. However, predators also tend to do quite well, and so strategies must (at least sometimes) respond to others' defections by defecting themselves. If a mechanism of "salary" or "tribute" can be developed (where strategies may give some of their fitness to other strategies), it might be possible to evolve "protector" strategies that play TIT FOR TAT. Such a situation seems to have some advantages: since there are only a few TIT FOR TAT players and many COOPERATE (or a similar variant) players, both TIT FOR TAT and COOPERATE will do well. After all, since COOPERATE is in the majority, TIT FOR TAT will rarely encounter other TIT FOR TAT players, and so noise will rarely cause strings of mutual defections. Meanwhile, predators will do poorly whenever they encounter TIT FOR TAT (the "police"). This sort of situation is what develops in most societies: we find that allowing neighbors to retaliate against one another has bad consequences, and so we establish a limited number of police to carry out any necessary retaliation. It would be quite interesting to model this as a Prisoner's Dilemma tournament, and to determine what conditions are necessary for the spontaneous development of such a "police" system.

REFERENCES

- Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books, Inc.
- Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. *Genetic algorithms and simulated annealing* (L. Davis, ed.). London: Pitman, 1987.
- Axelrod, R. (1997). *The complexity of cooperation*. New Jersey: Princeton University Press.
- Darwin, C. (1859). *On the origin of species*.
- Hay, I. (1916). *The first hundred thousand*. London: Wm. Blackwood.
P and the evolution of pathogen virulence. Zurich, Institute of Plant Sciences.
- Martino, C. (1995). *Emergent nastiness in iterated prisoner's dilemma games*. Massachusetts: MIT 2.725 Final Project.
- Sandholm, T, Crites, R. *Multiagent reinforcement learning in the iterated prisoner's dilemma*. Amherst: University of Massachusetts.
- Savage, D.C. (1977). Interactions between the Host and its Microbes., *Microbial ecology of the gut* (R.T. J. Clarke and T. Bauchop, eds.) 277-310. New York: Academic Press.
- Tyler, M. (2000). *Developmental biology, a guide for experimental study*. Massachusetts, Sinauer Associates, Inc.