

Evolving New Strategies

The Evolution of Strategies in the
Iterated Prisoner's Dilemma

Robert Axelrod

Adapted from Robert Axelrod, "The Evolution of Strategies in the Iterated Prisoner's Dilemma," in Lawrence Davis (ed.), Genetic Algorithms and Simulated Annealing, (London: Pitman, and Los Altos, CA: Morgan Kaufman, 1987), pp. 32-41.

In complex environments, individuals are not fully able to analyze the situation and calculate their optimal strategy.¹ Instead they can be expected to adapt their strategy over time based upon what has been effective and what has not. One useful analogy to the adaptation process is biological evolution. In evolution, strategies that have been relatively effective in a population become more widespread, and strategies that have been less effective become less common in the population.

Biological evolution has been highly successful at discovering complex and effective methods of adapting to very rich environmental situations. This is accomplished by differential reproduction of the more successful individuals. The evolutionary process also requires that successful characteristics be inherited through a genetic mechanism that allows some chance for new strategies to be discovered. One genetic mechanism allowing new strategies to be discovered is mutation. Another mechanism is crossover, whereby sexual reproduction takes some genetic material from one parent and some from the other.

The mechanisms that have allowed biological evolution to be so good at adaptation have been employed in the field of artificial intelligence. The artificial intelligence technique is called the "genetic algorithm" (Holland, 1975). While other methods of

representing strategies in games as finite automata have been used (Rubinstein 1986, Megiddo and Wigderson 1986, Miller 1989, Binmore and Samuelson 1990, Lomborg 1991), the genetic algorithm itself has not previously been used in game theoretic settings.

This paper will first demonstrate the genetic algorithm in the context of a rich social setting, the environment formed by the strategies submitted to a Prisoner's Dilemma computer tournament. The results show that the genetic algorithm is surprisingly successful at discovering complex and effective strategies that are well adapted to this complex environment. Next the paper shows how the results of this simulation experiment can be used to illuminate important issues in the evolutionary approach to adaptation, such as the relative advantage of developing new strategies based upon one or two parent strategies, the role of early commitments in the shaping of evolutionary paths, and the extent to which evolutionary processes are optimal or arbitrary.

The simulation method involves the following steps:

1. the specification of an environment in which the evolutionary process can operate,
2. the specification of the genetics, including the way in which information on the simulated chromosome is translated into a strategy for the simulated individual,
3. the design of an experiment to study the effects of alternative realities (such as repeating the experiment under identical conditions to see if random mutations lead to convergent or divergent evolutionary outcomes), and
4. the running of the experiment for a specified number of generations on a computer, and the statistical analysis of the results.

The Simulated Environment

An interesting set of environmental challenges are provided by the fact that many of the benefits sought by living things such as people are disproportionately available to cooperating groups. The problem is that while an individual can benefit from mutual cooperation, each one can also do even better by exploiting the cooperative efforts of others. Over a period of time, the same individuals may interact again, allowing for complex patterns of strategic interactions (Axelrod and Hamilton, 1981).

The Prisoner's Dilemma is an elegant embodiment of the problem of achieving mutual cooperation, and therefore provides the basis for the analysis. In the Prisoner's Dilemma, two individuals can each either cooperate or defect. The payoff to a player affects its reproductive success. No matter what the other does, the selfish choice of defection yields a higher payoff than cooperation. But if both defect, both do worse than if both had cooperated. Table 1-1 shows the payoff matrix of the Prisoner's Dilemma used in this study.

Table 1-1 here.

In many settings, the same two individuals may meet more than once. If an individual can recognize a previous interactant and remember some aspects of the prior outcomes, then the strategic situation becomes an iterated Prisoner's Dilemma. A strategy would take the form of a decision rule which specified the probability of cooperation or defection as a function of the history of the interaction so far.

To see what type of strategy can thrive in a variegated environment of more or less sophisticated strategies, I conducted a computer tournament for the Prisoner's Dilemma. The strategies were submitted by game theorists in economics sociology, political science, and mathematics (Axelrod, 1980a). The 14 entries and a totally random strategy were paired with each other in a round robin tournament. Some of the strategies were quite intricate. An example is one which on each move models the behavior of the

other player as a Markov process, and then uses Bayesian inference to select what seems the best choice for the long run. However, the result of the tournament was that the highest average score was attained by the simplest of all strategies, TIT FOR TAT. This strategy is simply one of cooperating on the first move and then doing whatever the other player did on the preceding move. Thus TIT FOR TAT is a strategy of cooperation based upon reciprocity.

The results of the first round were circulated and entries for a second round were solicited. This time there were 62 entries from six countries (Axelrod, 1980b). Most of the contestants were computer hobbyists, but there were also professors of evolutionary biology, physics, and computer science, as well as the five disciplines represented in the first round. TIT FOR TAT was again submitted by the winner of the first round, Anatol Rapoport. It won again.

The second round of the computer tournament provides an rich environment in which to test the evolution of behavior. It turns out that just eight of the entries can be used to account for how well a given rule did with the entire set. These eight rules can be thought of as representatives of the full set in the sense that the scores a given rule gets with them can be used to predict the average score the rule gets over the full set. In fact, 98% of the variance in the tournament scores is explained by knowing a rule's performance with these eight representatives. So these representative strategies can be used as a complex environment in which to evaluate an evolutionary simulation. What is needed next is a way of representing the genetic material of a population so that the evolutionary process can be studied in detail.

The Genetic Algorithm

The inspiration for how to conduct simulation experiments of genetics and evolution comes from an artificial intelligence procedure developed by computer scientist

John Holland and called the genetic algorithm (Holland 1975, 1980; Goldberg, 1989). For an excellent introduction to the genetic algorithm see Holland (1992) and Riolo (1992). The idea is based on the way in which a chromosome serves a dual purpose: it provides a representation of what the organism will become, and it also provides the actual material which can be transformed to yield new genetic material for the next generation.

Before going into details, it may help to give a brief overview of how the genetic algorithm works. The first step is to specify a way of representing each allowable strategy as a string of genes on a chromosome which can undergo genetic transformations, such as mutation. Then the initial population is constructed from the allowable set (perhaps by simply picking at random). In each generation, the effectiveness of each individual in the population is determined by running the individual in the current strategic environment. Finally, the relatively successful strategies are used to produce offspring which resemble the parents. Pairs of successful offspring are selected to mate and produce the offspring for the next generation. Each offspring draws part of its genetic material from one parent and part from another. Moreover, completely new material is occasionally introduced through mutation. After many generations of selection for relatively successful strategies, the result might well be a population that is substantially more successful in the given strategic environment than the original population.

To explain how the genetic algorithm can work in a game context, consider the strategies available for playing the iterated Prisoner's Dilemma. To be more specific, consider the set of strategies that are deterministic and use the outcomes of the three previous moves to make a choice in the current move. Since there are four possible outcomes for each move, there are $4 \times 4 \times 4 = 64$ different histories of the three previous moves. Therefore to determine its choice of cooperation or defection, a strategy would only need to determine what to do in each of the situations which could arise. This could be specified by a list of sixty-four C's and D's (C for cooperation and D for defection). For example, one of these sixty-four genes indicates whether the individual cooperates or

defects when in a rut of three mutual defections. Other parts of the chromosome would cover all the other situations that could arise.

To get the strategy started at the beginning of the game, it is also necessary to specify its initial premises about the three hypothetical moves which preceded the start of the game. To do this requires six more genes, making a total of seventy loci on the chromosome.² This string of seventy C's and D's would specify what the individual would do in every possible circumstance and would therefore completely define a particular strategy. The string of 70 genes would also serve as the individual's chromosome for use in reproduction and mutation.

There is a huge number of strategies which can be represented in this way. In fact, the number is 2 to the 70th power, which is about 10 to the 21st power.³ An exhaustive search for good strategies in this huge collection of strategies is clearly out of the question. If a computer had examined these strategies at the rate of 100 per second since the beginning of the universe, less than one percent would have been checked by now.

To find effective strategies in such a huge set, a very powerful technique is needed. This is where Holland's "genetic algorithm" comes in. It was originally inspired by biological genetics, but was adapted as a general problem solving technique. In the present context, it can be regarded as a model of a "minimal genetics" which can be used to explore theoretical aspects of evolution in rich environments. The outline of the simulation program works in five stages. See Table 1-2.

Table 1-2 here.

1. An initial population is chosen. In the present context the initial individuals can be represented by random strings of seventy C's and D's.

2. Each individual is run in the current environment to determine its effectiveness. In the present context this means that each individual player uses the strategy defined by

its chromosome to play an iterated Prisoner's Dilemma with other strategies, and the individual's score is its average over all the games it plays.⁴

3. The relatively successful individuals are selected to have more offspring. The method used is to give an average individual one mating, and to give two matings to an individual who is one standard deviation more effective than the average. An individual who is one standard deviation below the population average would then get no matings.

4. The successful individuals are then randomly paired off to produce two offspring per mating. For convenience, a constant population size is maintained. The strategy of an offspring is determined from the strategies of the two parents. This is done by using two genetic operators: crossover and mutation.

a. Crossover is a way of constructing the chromosomes of the two offspring from the chromosomes of two parents. It can be illustrated by an example of two parents, one of whom has seventy C's in its chromosome (indicating that it will cooperate in each possible situation that can arise), and the other of whom has seventy D's in its chromosome (indicating that it will always defect). Crossover selects one or more places to break the parents' chromosomes in order to construct two offspring each of whom has some genetic material from both parents. In the example, if a single break occurs after the third gene, then one offspring will have three C's followed by sixty-seven D's, while the other offspring will have three D's followed by sixty-seven D's.

b. Mutation in the offspring occurs by randomly changing a very small proportion of the C's to D's or visa versa.

5. This gives a new population. This new population will display patterns of behavior that are more like those of the successful individuals of the previous generation, and less like those of the unsuccessful ones. With each new generation, the individuals with relatively high scores will be more likely to pass on parts of their strategies, while the relatively unsuccessful individuals will be less likely to have any parts of their strategies passed on.

Simulation Results

The computer simulations were done using a population size of twenty individuals per generation. Levels of crossover and mutation were chosen averaging one crossover and one-half mutation per chromosome per generation. Each game consisted of 151 moves, the average game length used in the tournament. With each of the twenty individuals meeting eight representatives, this made for about 24,000 moves per generation. A run consisted of 50 generations. Forty runs were conducted under identical conditions to allow an assessment of the variability of the results.

The results are quite remarkable: from a strictly random start, the genetic algorithm evolved populations whose median member was just as successful as the best rule in the tournament, TIT FOR TAT. Most of the strategies that evolved in the simulation actually resemble TIT FOR TAT, having many of the properties that make TIT FOR TAT so successful. For example, five behavioral alleles in the chromosomes evolved in the vast majority of the individuals to give them behavioral patterns that were adaptive in this environment and mirrored what TIT FOR TAT would do in similar circumstances. These patterns are:

1. Don't rock the boat: continue to cooperate after three mutual cooperations (which can be abbreviated as C after RRR).
2. Be provokable: defect when the other player defects out of the blue (D after receiving RRS).
3. Accept an apology: continue to cooperate after cooperation has been restored (C after TSR).
4. Forget: cooperate when mutual cooperation has been restored after an exploitation (C after SRR).
5. Accept a rut: defect after three mutual defections (D after PPP).

The evolved rules behave with specific representatives in much the same way as TIT FOR TAT does. They did about as well as TIT FOR TAT did with each of the eight representatives. Just as TIT FOR TAT did, most of the evolved rules did well by achieving almost complete mutual cooperation with seven of the eight representatives. Like TIT FOR TAT, most of the evolved rules do poorly only with one representative, called ADJUSTER, that adjusts its rate of defection to try to exploit the other player. In all, 95% of the time the evolved rules make the same choice as TIT FOR TAT would make in the same situation.

While most of the runs evolve populations whose rules are very similar to TIT FOR TAT, in eleven of the forty runs, the median rule actually does substantially better than TIT FOR TAT.⁵ In these eleven runs, the populations evolved strategies that manage to exploit one of the eight representatives at the cost of achieving somewhat less cooperation with two others. But the net effect is a gain in effectiveness.

This is a remarkable achievement because to be able to get this added effectiveness, a rule must be able to do three things. First, it must be able to discriminate between one representative and another based upon only the behavior the other player shows spontaneously or is provoked into showing. Second, it must be able to adjust its own behavior to exploit a representative that is identified as an exploitable player. Third, and perhaps most difficult, it must be able to achieve this discrimination and exploitation without getting into too much trouble with the other representatives. This is something that none of the rules originally submitted to the tournament were able to do.

These very effective rules evolved by breaking the most important device developed in the computer tournament, namely to be "nice", that is never to be the first to defect. These highly effective rules always defect on the very first move, and sometimes on the second move as well, and use the choices of the other player to discriminate what should be done next. The highly effective rules then had responses that allowed them to "apologize" and get to mutual cooperation with most of the unexploitable representatives,

and they had different responses which allowed them to exploit a representative that was exploitable.

While these rules are highly effective, it would not accurate to say that they are better than TIT FOR TAT. While they are better in the particular environment consisting of fixed proportions of the eight representatives of the second round of the computer tournament, they are probably not very robust in other environments. Moreover, in an ecological simulation these rules would be destroying the basis of their own success as the exploited representative would become a smaller and smaller part of the environment (Axelrod 1984, pp. 49-52 and 203-5). While the genetic algorithm was sometimes able to evolve rules that are more effective than any entry in the tournament, the algorithm was only able to do so by trying many individuals in many generations against a fixed environment. In sum, the genetic algorithm is very good at what actual evolution does so well: developing highly specialized adaptations to specific environmental settings.

In the evolution of these highly effective strategies, the computer simulation employed sexual reproduction, where two parents contributed genetic material to each offspring. To see what would happen with asexual reproduction, forty additional runs were conducted in which only one parent contributed genetic material to each offspring. In these runs, the populations still evolved toward rules that did about as well as TIT FOR TAT in most cases. However, the asexual runs were only half as likely to evolve populations in which the median member was substantially more effective than TIT FOR TAT.⁶

So far, the simulation experiments have dealt with populations evolving in the context of a constant environment. What would happen if the environment is also changing? To examine this situation, another simulation experiment with sexual reproduction was conducted in which the environment consisted of the evolving population itself. In this experiment each individual plays the iterated Prisoner's

Dilemma with each other member of the population rather than play with the eight representatives. At any given time, the environment can be quite complex. For an individual to do well requires that its strategy achieves a high average effectiveness with the nineteen other strategies that are also present in the population. Thus as the more effective rules have more offspring, the environment itself changes. In this case, adaptation must be done in the face of a moving target. Moreover, the selection process is frequency dependent, meaning that the effectiveness of a strategy depends upon what strategies are being used by the other members of the population.

The results of the ten runs conducted in this manner display a very interesting pattern. For a typical run, see Figure 1-1. From a random start, the population evolves away from whatever cooperation was initially displayed. The less cooperative rules do better than the more cooperative rules because at first there are few other players who are responsive - and when the other player is unresponsive the most effective thing for an individual to do is simply defect. This decreased cooperation in turn causes everyone to get lower scores as mutual defection becomes more and more common. However, after about ten or twenty generations the trend starts to reverse. Some players evolve a pattern of reciprocating what cooperation they find, and these reciprocating players tend to do well because they can do very well with others who reciprocate without being exploited for very long by those who just defect. The average scores of the population then start to increase as cooperation based upon reciprocity becomes better and better established. So the evolving social environment led to a pattern of decreased cooperation and decreased effectiveness, followed by a complete reversal based upon an evolved ability to discriminate between those who will reciprocate cooperation and those who won't. As the reciprocators do well, they spread in the population resulting in more and more cooperation and greater and greater effectiveness.

Figure 1-1 here.

Conclusions

1. The genetic algorithm is a highly effective method of searching for effective strategies in a huge space of possibilities. Following Quincy Wright (1977, pp. 452-454), the problem for evolution can be conceptualized as a search for relatively high points in a multidimensional field of gene combinations, where height corresponds to fitness. When the field has many local optima, the search becomes quite difficult. When the number of dimensions in the field becomes great, the search is even more difficult. What the computer simulations demonstrate is that minimal system of the genetic algorithm is a highly efficient method for searching such a complex multidimensional space. The first experiment shows that even with a seventy dimensional field of genes, quite effective strategies can be found within fifty generations. Sometimes the genetic algorithm found combinations of genes that violate the previously accepted mode of operation (not being the first to defect) to achieve even greater effectiveness than had been thought possible.

2. Sexual reproduction does indeed help the search process. This was demonstrated by the much increased chance of achieving highly effective populations in the sexual experiment compared to the asexual experiment.⁷

3. Some aspects of evolution are arbitrary. In natural settings, one might observe that a population has little variability in a specific gene. In other words one of the alleles for that gene has become fixed throughout the population. One might be tempted to assume from this that the allele is more adaptive than any alternative allele. However, this may not be the case. The simulation of evolution allows an exploration of this possibility by allowing repetitions of the same conditions to see just how much variability there is in the outcomes. In fact, the simulations show two reasons why convergence in a population may actually be arbitrary.

a. Genes that do not have much effect on the fitness of the individual may become fixed in a population because they "hitch-hike" on other genes that do (Maynard Smith and Haigh, 1974). For example, in the simulations some sequences of three moves may very rarely occur, so what the corresponding genes dictate in these situations may not matter very much. However, if the entire population are descendants of just a few individuals, then these irrelevant genes may be fixed to the values that their ancestors happened to share. Repeated runs of a simulation allow one to notice that some genes become fixed in one population but not another, or that they become fixed in different ways in different populations.

b. In some cases, some parts of the chromosome are arbitrary in content, but what is not arbitrary is that they be held constant. By being fixed, other parts of the chromosome can adapt to them. For example, the simulations of the individual chromosomes had six genes devoted to coding for the premises about the three moves that preceded the first move in the game. When the environment was the eight representatives, the populations in different runs of the simulation developed different premises. Within each run, however, the populations were usually very consistent about the premises: the six premise genes had become fixed. Moreover, within each population these genes usually became fixed quite early. It is interesting that different populations evolved quite different premises. What was important for the evolutionary process was to fix the premise about which history is assumed at the start so that the other parts of the chromosome could adapt on the basis of a given premise.

4. There is a tradeoff between the gains to be made from flexibility and the gains to be made from commitment and specialization (March, 1991). Flexibility might help in the long run, but in an evolutionary system, the individuals also have to survive in the short run if they are to reproduce. This feature of evolution arises at several levels.

a. As the simulations have shown, the premises became fixed quite early. This meant a commitment to which parts of the chromosome would be consulted in the

first few moves, and this in turn meant giving up flexibility as more and more of the chromosome evolved on the basis of what had been fixed. This in turn meant that it would be difficult for a population to switch to a different premise. So flexibility was given up so that the advantages of commitment could be reaped.

b. There is also a tradeoff between short and long term gains in the way selection was done in the simulation experiments. In any given generation there would typically be some individuals that did much better than the average, and some that did only a little better than the average. In the short run, the way to maximize the expected performance of the next generation would be to have virtually all of the offspring come from the very best individuals in the present generation. But this would imply a rapid reduction in the genetic variability of the population, and a consequent slowing of the evolutionary process later on. If the moderately successful were also given a chance to have some offspring, this would help the long term prospects of the population at the cost of optimizing in the short run. Thus there is an inherent tradeoff between exploitation and exploration, i.e. between exploiting what already works best and exploring possibilities that might eventually evolve into something even better (Holland, 1975, p. 160).

5. Evolutionary commitments can be irreversible. For example, in most of the populations facing the environment of the eight representatives, the individuals evolved strategies which are very similar to TIT FOR TAT. Since TIT FOR TAT had done best in the computer tournament itself, I did not think that it would be possible to do much better with an evolutionary process. But as noted earlier, in about a quarter of the simulation runs with sexual reproduction, the population did indeed evolve substantially better strategies - strategies which were quite different from TIT FOR TAT. These highly effective strategies defected on the very first move, and often on the second move as well, in order get information to determine whether the other player was the type that could be exploited or not. The more common populations of strategies cooperated from

the beginning and employed reciprocity in a manner akin to TIT FOR TAT. While these more common strategies might easily mutate to try a defection at the start of the game, such behavior would be extremely costly unless the individual already had effective ways of using the information that resulted. Moreover, once the population had evolved to be about as effective as TIT FOR TAT, such mutation would have to be quite effective in order to survive long enough to be perfected. Thus, once the population takes a specific route (in this case towards reciprocity) it can easily become trapped in a local maxima. Indeed, only the fact that enough simulation runs were conducted lead to the discovery that in this particular environment reciprocity was only a local maxima, and that something better was in fact possible. In a field situation such a discovery might not be possible since there might be essentially just one gene pool.

Topics Amenable to Simulation

The methodology for the genetic simulation developed in this paper can be used to explore learning processes in game theoretic settings. Here is a list of issues that can be studied with genetic simulations, inspired by analogs to evolutionary biology.

1. Mutation. The simulation approach developed here suggests that there is an inherent tradeoff for a gene pool between exploration of possibilities (best done with a high mutation rate), and exploitation of the possibilities already contained in the current gene pool (best done with a low mutation rate). This in turn suggests the advantage of having mutation rates adapt to the rate of change in the environment.⁸

2. Crossover. In sexual reproduction, crossover serves to give each offspring genetic material from both parents. Crossover rates that are too low would frequently give whole chromosomes of genetic material from a single parent to an offspring. But crossover rates that are too high would frequently split up co-adapted sets of alleles that are on the same chromosome. Perhaps the existence of a multiplicity of chromosomes

(rather than one long chromosome) is more than a mechanical convenience, but is an adaptation to the need from low crossover rates without the disadvantage of having each offspring being likely to get genetic material from only one parent.

3. Inversion. Inversion changes the order of the genes in a chromosome. It can bring co-adapted sets of alleles closer together on the chromosome so that they will be split apart by crossover less often. How is the ideal rate of inversion determined?

4. Coding principles. Biological chromosomes are known to contain material that does not directly code for proteins, but performs other roles such as marking the boundaries of genes, or perhaps provides no function at all. Genetic material may also appear in highly redundant form in the chromosome. Genetic simulation experiments might shed new light on the theoretical implications of various coding schemes, and their possible role in error reduction and regulation. Or they might show how some genetic material can exist as "free riders."

5. Dominant and recessive genes. Mendel's famous experiments demonstrate that dominant and recessive alleles serve to overcome Darwin's concern that blending of parental characteristics would eliminate the variability of a population. Genetic simulation can be used to explore the implications of these and other genetic mechanisms for the maintenance of population variability in the face of selection pressure for local optimality. In particular, it should be possible to explore just which types of phenotypic features are best coded in terms of dominant and recessive genes, and which are best coded in other systems of genetic expression.

6. Gradual vs. punctuated evolution. Genetic simulation experiments might also shed light on the contemporary debate about whether evolution proceeds in gradual steps or whether it tends to move it fits and starts. This type of work might require simulations of tens of thousands of generations, but runs of such length are feasible.

7. Population viscosity. Obstacles to random mating may exist due to geographic or other forces tending to favor subdivisions of the population. Some computer modeling

has already been done for models of this type (Boorman and Levitt, 1980, pp. 78-87; Tanese 1989) revealing clues about the qualitative features of the spread of a social trait based upon frequency dependent selection.

8. Speciation and ecological niches. When distinct ecological niches exist, a single species tends to differentiate into two or more species to take advantage of the different opportunities offered by the different niches. In learning terms, differentiation into two or more species means that a new strategy is formed from ideas represented in only part of the total population. Genetic simulation can explore this process by helping to specify the conditions under which the advantages of specialization outweigh the disadvantages of narrower mating opportunities and reduced ecological flexibility. The fundamental point is that thinking about genetics as a simulation problem gives a new perspective on the functioning of learning processes.

The genetic simulations provided in this paper are highly abstract systems. The populations are very small, and the number of generations is few. More significantly, the genetic process have only two operators, mutation and crossover. Compared to biological genetics, this is a highly simplified system. Nevertheless, the genetic algorithm displayed a remarkable ability to evolve sophisticated and effective strategies in a complex environment.

References

- Axelrod, Robert, 1980a. "Effective Choice in the Prisoner's Dilemma," Journal of Conflict Resolution 24: 3-25.
- _____, 1980b. "More Effective Choice in the Prisoner's Dilemma," Journal of Conflict Resolution 24: 379-403.
- _____, 1984. The Evolution of Cooperation. New York: Basic Books.
- _____, and William D. Hamilton, 1981. "The Evolution of Cooperation," Science 211:1390-96.
- Binmore, Ken and Larry Samuelson, 1990. "Evolutionary Stability in Repeated Games Played by Finite Automata," Center for Research on Economic and Social Theory, Working Paper 90-17, Department of Economics, University of Michigan, Ann Arbor, MI.
- Boorman, Scott A. and Paul R. Levitt, 1980. The Genetics of Altruism. New York: Academic Press.
- Goldberg, D. E. 1983. Computer-aided Gas Pipeline Operation Using Genetic Algorithms and Machine Learning. Ph.D. Dissertation University of Michigan (Civil Engineering).
- _____, 1989. Genetic Algorithms in Search, Optimization, and Machine Learning. Reading, MA: Addison-Wesley.
- Grefenstette, John, J. (ed.), 1985. Proceedings of an International Conference on Genetic Algorithms and Their Applications. Pittsburgh, PA: The Robotics Institute of Carnegie-Mellon University.
- Lomborg, Bjorn, 1991. "An Evolution of Cooperation," Masters Thesis, Institute of Political Science, University of Aarhus, Denmark.
- Hamilton, William D., 1980. "Sex versus Non-Sex versus Parasite," Oikos 35:282-290.

- _____, 1982. "Heritable True Fitness and Bright Birds: A Role for Parasites," Science 218: 384-387.
- _____, Robert Axelrod and Reiko Tanese, 1990. "Sexual Reproduction as an Adaptation to Resist Parasites." Proceedings of the National Academy of Sciences USA 87: 3566-3573.
- Holland, John H., 1975. Adaptation in Natural and Artificial Systems. Ann Arbor: University of Michigan Press.
- _____, 1980. "Adaptive Algorithms for Discovering and Using General Patterns in Growing Knowledge Bases," International Journal of Policy Analysis and Information Systems 4: 245-268.
- _____ 1992. "Genetic Algorithms," Scientific American 267:66-72, July.
- March, James G., 1991. "Exploration and Exploitation in Organizational Learning," Organizational Science, 2: 71-87.
- Maynard Smith, J. and J. Haigh, 1974. "The Hitch-hiking Effect of a Favorable Gene." Genet. Res., Camb. 23: 23-35.
- Megiddo, Nimrod, and Avi Wigderson, 1986. "On Play by Means of Computing Machines," IBM Research Division, BJ 4984 (52161), Yorktown Heights, New York.
- Miller, John, 1989. "The Coevolution of Automata in the Repeated Prisoner's Dilemma," Working Paper, 89-003, Santa Fe Institute, Santa Fe, NM.
- Tanese, Reiko, 1989. Distributed Genetic Algorithms for Function Optimization. Ph.D. Dissertation, University of Michigan (Computer Science and Engineering).
- Riolo, Rick L., 1992. "Survival of the Fittest Bits," Scientific American 267:114-16, July.
- Rubinstein, Ariel, 1986. "Finite Automata Play the Repeated Prisoner's Dilemma," Journal of Economic Theory, 39: 83-96.
- Wright, Sewall, 1977. Evolution and the Genetics of Populations, Volume 4, Experimental Results and Evolutionary Deductions. Chicago: University of

Chicago Press.

Table 1-2

The Basic Simulation

- * set up initial population with random chromosomes

- * for each of 50 generations
 - for each of 20 individuals
 - * for each of the 8 representatives
 - use premise part of the chromosome as individual's assumption about the three previous moves

 - for each of 151 moves
 - * make the individual's choice of cooperate (C) or defect (D) based upon the gene that encodes what to do given the three previous moves

 - * make the representative's choice of C or D based upon its own strategy applied to the history of the game so far

 - * update the individual's score based upon the outcome of this move (add 3 points if both


cooperated, add 5 points if the representative cooperated and the individual defected, etc.)

- reproduce the next generation

* for each individual assign the likely number of matings based upon the scaling function (1 for an average score, 2 for a score one standard deviation above average, etc.)

* for each of 10 matings construct two offspring from the two selected parents using crossover and mutation

**Figure 1-1
Prisoner's Dilemma
Evolving Environment**



¹I thank Stephanie Forrest and Reiko Tanese for their help with the computer programming, Michael D. Cohen and John Holland for their helpful suggestions, and the Harry Frank Guggenheim Foundation and the National Science Foundation for their financial support.

²The six premise genes encode the presumed C or D choices made by the individual and the other player in each of the three moves before the interaction actually begins.

³Some of these chromosomes give rise to equivalent strategies since certain genes might code for histories that could not arise given how loci are set. This does not necessarily make the search process any easier, however.

⁴The score is actually a weighted average of its scores with the eight representative, the weights having been chosen to give the best representation of the entire set of strategies in the second round of the tournament.

⁵The criterion for being substantially better than TIT FOR TAT is a median score of 450 points, which compares to TIT FOR TAT's weighted score of 428 with these eight representatives.

⁶This happened in 5 of the 40 runs with asexual reproduction compared to 11 of the 40 runs with sexual reproduction. This difference is significant at the .05 level using the one tailed chi-squared test.

⁷In biology, sexual reproduction comes at the cost of reduced fecundity. Thus if males provide little or no aid to offspring, a high (up to 2-fold) average extra fitness has to emerge as a property of sexual reproduction if sex is to be stable. The advantage must presumably come from recombination but has been hard to identify in biology. A simulation model has demonstrated that the advantage may well lie in the necessity to recombine defenses to defeat numerous parasites (Hamilton, Axelrod and Tanese, 1990). Unlike biology, in artificial intelligence applications, the added (computational) cost of sexuality is small.

⁸I owe this suggestion to Michael D. Cohen.