THE EVOLUTION OF COOPERATION

Robert Axelrod

Basic Books, Inc., Publishers *New York*

An earlier version of chapter 2, tables 1–5, and figures 1–2 appeared in *The Journal of Conflict Resolution* (1980).

An earlier version of chapter 3 appeared in the *American Political Science Review* (1981).

An earlier version of chapter 5 appeared in *Science* 211 (27 March 1981):1390–96.

Library of Congress Cataloging in Publication Data

Axelrod, Robert M.
The evolution of cooperation.

Bibliography: p. 223
Includes index.
1. Cooperativeness. 2. Games of strategy (Mathematics) 3. Conflict management. 4. Egoism. 5. Consensus (Social sciences) 6. Social interaction. I. Title.
HM131.A89 1984 302′.14 83-45255
ISBN 0-465-02122-0

Printed in the United States of America
Designed by Vincent Torre
10 9 8 7 6 5 4 3 2 1

CHAPTER 1

The Problem of Cooperation

UNDER WHAT CONDITIONS will cooperation emerge in a world of egoists without central authority? This question has intrigued people for a long time. And for good reason. We all know that people are not angels, and that they tend to look after themselves and their own first. Yet we also know that cooperation does occur and that our civilization is based upon it. But, in situations where each individual has an incentive to be selfish, how can cooperation ever develop?

The answer each of us gives to this question has a fundamental effect on how we think and act in our social, political, and economic relations with others. And the answers that others give have a great effect on how ready they will be to cooperate with us.

The most famous answer was given over three hundred years ago by Thomas Hobbes. It was pessimistic. He argued that before governments existed, the state of nature was dominated by the problem of selfish individuals who competed on such ruthless terms that life was "solitary, poor, nasty, brutish, and short" (Hobbes 1651/1962, p. 100). In his view, cooperation could not develop without a central authority, and consequently a strong government was necessary. Ever since, arguments about the proper scope of government have often focused on whether one could, or could not, expect cooperation to emerge in a particular domain if there were not an authority to police the situation.

Today nations interact without central authority. Therefore the requirements for the emergence of cooperation have relevance to many of the central issues of international politics. The most important problem is the security dilemma: nations often seek their own security through means which challenge the security of others. This problem arises in such areas as escalation of local conflicts and arms races. Related problems occur in international relations in the form of competition within alliances, tariff negotiations, and communal conflict in places like Cyprus.[1]

The Soviet invasion of Afghanistan in 1979 presented the United States with a typical dilemma of choice. If the United States continued business as usual, the Soviet Union might be encouraged to try other forms of noncooperative behavior later on. On the other hand, any substantial lessening of United States cooperation risked some form of retaliation, which could then set off counter-retaliation, setting up a pattern of mutual hostility that could be difficult to end. Much of the domestic debate about foreign policy is concerned with problems of just this type. And properly so, since these are hard choices.

In everyday life, we may ask ourselves how many times

we will invite acquaintances for dinner if they never invite us over in return. An executive in an organization does favors for another executive in order to get favors in exchange. A journalist who has received a leaked news story gives favorable coverage to the source in the hope that further leaks will be forthcoming. A business firm in an industry with only one other major company charges high prices with the expectation that the other firm will also maintain high prices—to their mutual advantage and at the expense of the consumer.

For me, a typical case of the emergence of cooperation is the development of patterns of behavior in a legislative body such as the United States Senate. Each senator has an incentive to appear effective to his or her constituents, even at the expense of conflicting with other senators who are trying to appear effective to *their* constituents. But this is hardly a situation of completely opposing interests, a zero-sum game. On the contrary, there are many opportunities for mutually rewarding activities by two senators. These mutually rewarding actions have led to the creation of an elaborate set of norms, or folkways, in the Senate. Among the most important of these is the norm of reciprocity—a folkway which involves helping out a colleague and getting repaid in kind. It includes vote trading but extends to so many types of mutually rewarding behavior that "it is not an exaggeration to say that reciprocity is a way of life in the Senate" (Matthews 1960, p. 100; see also Mayhew 1975).

Washington was not always like this. Early observers saw the members of the Washington community as quite unscrupulous, unreliable, and characterized by "falsehood, deceit, treachery" (Smith 1906, p. 190). In the 1980s the practice of reciprocity is well established. Even the significant changes in the Senate over the last two decades, tending toward more decentralization, more openness, and more equal distribution of power, have come without abating the folkway of reciprocity (Ornstein, Peabody, and Rhode 1977). As will be seen, it is *not* necessary to assume that senators are more honest, more generous, or more public-spirited than in earlier years to explain how cooperation based on reciprocity has emerged or proved stable. The emergence of cooperation can be explained as a consequence of individual senators pursuing their own interests.

The approach of this book is to investigate how individuals pursuing their own interests will act, followed by an analysis of what effects this will have for the system as a whole. Put another way, the approach is to make some assumptions about individual motives and then deduce consequences for the behavior of the entire system (Schelling 1978). The case of the U.S. Senate is a good example, but the same style of reasoning can be applied to other settings.

The object of this enterprise is to develop a theory of cooperation that can be used to discover what is necessary for cooperation to emerge. By understanding the conditions that allow it to emerge, appropriate actions can be taken to foster the development of cooperation in a specific setting.

The Cooperation Theory that is presented in this book is based upon an investigation of individuals who pursue their own self-interest without the aid of a central authority to force them to cooperate with each other. The reason for assuming self-interest is that it allows an examination of the difficult case in which cooperation is not completely based upon a concern for others or upon the welfare of the group as a whole. It must, however, be stressed that this assumption is actually much less restrictive than it appears.

If a sister is concerned for the welfare of her brother, the sister's self-interest can be thought of as including (among many other things) this concern for the welfare of her brother. But this does not necessarily eliminate all potential for conflict between sister and brother. Likewise a nation may act in part out of regard for the interests of its friends, but this regard does not mean that even friendly countries are always able to cooperate for their mutual benefit. So the assumption of self-interest is really just an assumption that concern for others does not completely solve the problem of when to cooperate with them and when not to.

A good example of the fundamental problem of cooperation is the case where two industrial nations have erected trade barriers to each other's exports. Because of the mutual advantages of free trade, both countries would be better off if these barriers were eliminated. But if either country were to unilaterally eliminate its barriers, it would find itself facing terms of trade that hurt its own economy. In fact, whatever one country does, the other country is better off retaining its own trade barriers. Therefore, the problem is that each country has an incentive to retain trade barriers, leading to a worse outcome than would have been possible had both countries cooperated with each other.

This basic problem occurs when the pursuit of self-interest by each leads to a poor outcome for all. To make headway in understanding the vast array of specific situations which have this property, a way is needed to represent what is common to these situations without becoming bogged down in the details unique to each. Fortunately, there is such a representation available: the famous *Prisoner's Dilemma* game.[2]

In the Prisoner's Dilemma game, there are two players. Each has two choices, namely cooperate or defect. Each must make the choice without knowing what the other will do. No matter what the other does, defection yields a higher payoff than cooperation. The dilemma is that if both defect, both do worse than if both had cooperated. This simple game will provide the basis for the entire analysis used in this book.

The way the game works is shown in figure 1. One player chooses a row, either cooperating or defecting. The other player simultaneously chooses a column, either cooperating or defecting. Together, these choices result in one of the four possible outcomes shown in that matrix. If both players cooperate, both do fairly well. Both get R, the *reward for mutual cooperation.* In the concrete illustration of figure 1 the reward is 3 points. This number might, for example, be a payoff in dollars that each player gets for that outcome. If one player cooperates but the other defects, the defecting player gets the *temptation to defect,* while the cooperating player gets the *sucker's payoff.* In the example, these are 5 points and 0 points respectively. If both defect, both get 1 point, the *punishment for mutual defection.*

What should you do in such a game? Suppose you are the row player, and you think the column player will cooperate. This means that you will get one of the two outcomes in the first column of figure 1. You have a choice. You can cooperate as well, getting the 3 points of the reward for mutual cooperation. Or you can defect, getting the 5 points of the temptation payoff. So it pays to defect if you think the other player will cooperate. But now suppose that you think the other player will defect. Now you are in the second column of figure 1, and you have a choice between cooperating, which would make you a sucker and give you 0 points, and defecting, which would result in, mutual punishment giving you 1 point. So it pays to defect if you think the other player will defect. This means that it is better to defect if you think the other player will cooperate, *and* it is better to defect if you think the other player will defect. So no matter what the other player does, it pays for you to defect.

So far, so good. But the same logic holds for the other player too. Therefore, the other player should defect no matter what you are expected to do. So you should both defect. But then you both get 1 point which is worse than the 3 points of the reward that you both could have gotten had you both cooperated. Individual rationality leads to a worse outcome for both than is possible. Hence the dilemma.

The Prisoner's Dilemma is simply an abstract formulation of some very common and very interesting situations in which what is best for each person individually leads to mutual defection, whereas everyone would have been better off with mutual cooperation. The definition of Prisoner's Dilemma requires that several relationships hold among the four different potential outcomes. The first relationship specifies the order of the four payoffs. The best a player can do is get T, the temptation to defect when the

FIGURE 1
The Prisoner's Dilemma

		Column Player	
		Cooperate	*Defect*
Row Player	Cooperate	R=3, R=3 Reward for mutual cooperation	S=0, T=5 Sucker's payoff, and temptation to defect
	Defect	T=5, S=0 Temptation to defect and sucker's payoff	P=1, P=1 Punishment for mutual defection

NOTE: The payoffs to the row chooser are listed first.

other player cooperates. The worst a player can do is get *S*, the sucker's payoff for cooperating while the other player defects. In ordering the other two outcomes, *R*, the reward for mutual cooperation, is assumed to be better than *P*, the punishment for mutual defection. This leads to a preference ranking of the four payoffs from best to worst as *T*, *R*, *P*, and *S*.

The second part of the definition of the Prisoner's Dilemma is that the players cannot get out of their dilemma by taking turns exploiting each other. This assumption means that an even chance of exploitation and being exploited is not as good an outcome for a player as mutual cooperation. It is therefore assumed that the reward for mutual cooperation is greater than the average of the temptation and the sucker's payoff. This assumption, together with the rank ordering of the four payoffs, defines the Prisoner's Dilemma.

Thus two egoists playing the game *once* will both choose their dominant choice, defection, and each will get less than they both could have gotten if they had cooperated. If the game is played a known finite number of times, the players still have no incentive to cooperate. This is certainly true on the last move since there is no future to influence. On the next-to-last move neither player will have an incentive to cooperate since they can both anticipate a defection by the other player on the very last move. Such a line of reasoning implies that the game will unravel all the way back to mutual defection on the first move of any sequence of plays that is of known finite length (Luce and Raiffa 1957, pp. 94–102). This reasoning does not apply if the players will interact an indefinite number of times. And in most realistic settings, the players cannot be sure when the last interaction between them will take place. As will be

shown later, with an indefinite number of interactions, cooperation can emerge. The issue then becomes the discovery of the precise conditions that are necessary and sufficient for cooperation to emerge.

In this book I will examine interactions between just two players at a time. A single player may be interacting with many others, but the player is assumed to be interacting with them one at a time.[3] The player is also assumed to recognize another player and to remember how the two of them have interacted so far. This ability to recognize and remember allows the history of the particular interaction to be taken into account by a player's strategy.

A variety of ways to resolve the Prisoner's Dilemma have been developed. Each involves allowing some additional activity that alters the strategic interaction in such a way as to fundamentally change the nature of the problem. The original problem remains, however, because there are many situations in which these remedies are not available. Therefore, the problem will be considered in its fundamental form, without these alterations.

1. There is no mechanism available to the players to make enforceable threats or commitments (Schelling 1960). Since the players cannot commit themselves to a particular strategy, each must take into account all possible strategies that might be used by the other player. Moreover the players have all possible strategies available to themselves.

2. There is no way to be sure what the other player will do on a given move. This eliminates the possibility of metagame analysis (Howard 1971) which allows such options as "make the same choice as the other is about to make." It also eliminates the possibility of reliable reputations such as might be based on watching the other player interact with

third parties. Thus the only information available to the players about each other is the history of their interaction so far.

3. There is no way to eliminate the other player or run away from the interaction. Therefore each player retains the ability to cooperate or defect on each move.

4. There is no way to change the other player's payoffs. The payoffs already include whatever consideration each player has for the interests of the other (Taylor 1976, pp. 69–73).

Under these conditions, words not backed by actions are so cheap as to be meaningless. The players can communicate with each other only through the sequence of their own behavior. This is the problem of the Prisoner's Dilemma in its fundamental form.

What makes it possible for cooperation to emerge is the fact that the players might meet again. This possibility means that the choices made today not only determine the outcome of this move, but can also influence the later choices of the players. The future can therefore cast a shadow back upon the present and thereby affect the current strategic situation.

But the future is less important than the present—for two reasons. The first is that players tend to value payoffs less as the time of their obtainment recedes into the future. The second is that there is always some chance that the players will not meet again. An ongoing relationship may end when one or the other player moves away, changes jobs, dies, or goes bankrupt.

For these reasons, the payoff of the next move always counts less than the payoff of the current move. A natural way to take this into account is to cumulate payoffs over time in such a way that the next move is worth some frac-

tion of the current move (Shubik 1970). The *weight* (or importance) of the next move relative to the current move will be called w. It represents the degree to which the payoff of each move is discounted relative to the previous move, and is therefore a *discount parameter.*

The discount parameter can be used to determine the payoff for a whole sequence. To take a simple example, suppose that each move is only half as important as the previous move, making $w = 1/2$. Then a whole string of mutual defections worth one point each move would have a value of 1 on the first move, $1/2$ on the second move, $1/4$ on the third move, and so on. The cumulative value of the sequence would be $1 + 1/2 + 1/4 + 1/8 \ldots$ which would sum to exactly 2. In general, getting one point on each move would be worth $1 + w + w^2 + w^3 \ldots$. A very useful fact is that the sum of this infinite series for any w greater than zero and less than one is simply $1/(1-w)$. To take another case, if each move is worth 90 percent of the previous move, a string of 1's would be worth ten points because $1/(1-w) = 1/(1-.9) = 1/.1 = 10$. Similarly, with w still equal to .9, a string of 3 point mutual rewards would be worth three times this, or 30 points.

Now consider an example of two players interacting. Suppose one player is following the policy of always defecting (ALL D), and the other player is following the policy of TIT FOR TAT. TIT FOR TAT is the policy of cooperating on the first move and then doing whatever the other player did on the previous move. This policy means that TIT FOR TAT will defect once after each defection of the other player. When the other player is using TIT FOR TAT, a player who always defects will get T on the first move, and P on all subsequent moves. The *value* (or *score*) to someone using ALL D when playing with someone using TIT FOR TAT is thus the sum of T for the first move, wP for the second move, w^2P for the third move, and so on.[4]

Both ALL D and TIT FOR TAT are strategies. In general, a *strategy* (or *decision rule*) is a specification of what to do in any situation that might arise. The situation itself depends upon the history of the game so far. Therefore, a strategy might cooperate after some patterns of interaction and defect after others. Moreover, a strategy may use probabilities, as in the example of a rule which is entirely random with equal probabilities of cooperation and defection on each move. A strategy can also be quite sophisticated in its use of the pattern of outcomes in the game so far to determine what to do next. An example is one which, on each move, models the behavior of the other player using a complex procedure (such as a Markov process), and then uses a fancy method of statistical inference (such as Bayesian analysis) to select what seems the best choice for the long run. Or it may be some intricate combination of other strategies.

The first question you are tempted to ask is, "What is the best strategy?" In other words, what strategy will yield a player the highest possible score? This is a good question, but as will be shown later, no best rule exists independently of the strategy being used by the other player. In this sense, the iterated Prisoner's Dilemma is completely different from a game like chess. A chess master can safely use the assumption that the other player will make the most feared move. This assumption provides a basis for planning in a game like chess, where the interests of the players are completely antagonistic. But the situations represented by the Prisoner's Dilemma game are quite different. The interests of the players are not in total conflict. Both players can do

well by getting the reward, R, for mutual cooperation or both can do poorly by getting the punishment, P, for mutual defection. Using the assumption that the other player will always make the move you fear most will lead you to expect that the other will never cooperate, which in turn will lead you to defect, causing unending punishment. So unlike chess, in the Prisoner's Dilemma it is not safe to assume that the other player is out to get you.

In fact, in the Prisoner's Dilemma, the strategy that works best depends directly on what strategy the other player is using and, in particular, on whether this strategy leaves room for the development of mutual cooperation. This principle is based on the weight of the next move relative to the current move being sufficiently large to make the future important. In other words, the discount parameter, w, must be large enough to make the future loom large in the calculation of total payoffs. After all, if you are unlikely to meet the other person again, or if you care little about future payoffs, then you might as well defect now and not worry about the consequences for the future.

This leads to the first formal proposition. It is the sad news that if the future is important, there is no one best strategy.

Proposition 1. If the discount parameter, w, is sufficiently high, there is no best strategy independent of the strategy used by the other player.

The proof itself is not hard. Suppose that the other player is using ALL D, the strategy of always defecting. If the other player will never cooperate, the best you can do is always to defect yourself. Now suppose, on the other hand, that the other player is using a strategy of "permanent retaliation." This is the strategy of cooperating until you defect and then always defecting after that. In that case, your best strategy is never to defect, provided that the temptation to defect on the first move will eventually be more than compensated for by the long-term disadvantage of getting nothing but the punishment, *P*, rather than the reward, *R*, on future moves. This will be true whenever the discount parameter, w, is sufficiently great.[5] Thus, whether or not you should cooperate, even on the first move, depends on the strategy being used by the other player. Therefore, if w is sufficiently large, there is no one best strategy.

In the case of a legislature such as the U.S. Senate, this proposition says that if there is a large enough chance that a member of the legislature will interact *again* with another member, there is no one best strategy to use independently of the strategy being used by the other person. It would be best to cooperate with someone who will reciprocate that cooperation in the future, but not with someone whose future behavior will not be very much affected by this interaction (see, for example, Hinckley 1972). The very possibility of achieving stable mutual cooperation depends upon there being a good chance of a continuing interaction, as measured by the magnitude of w. As it happens, in the case of Congress, the chance of two members having a continuing interaction has increased dramatically as the biennial turnover rates have fallen from about 40 percent in the first forty years of the republic to about 20 percent or less in recent years (Young 1966, pp. 87–90; Polsby 1968; Jones 1977, p. 154; Patterson 1978, pp. 143–44).

However, saying that a continuing chance of interaction is necessary for the development of cooperation is not the same as saying that it is sufficient. The demonstration that there is not a single best strategy leaves open the question

of what patterns of behavior can be expected to emerge when there actually is a sufficiently high probability of continuing interaction between two individuals.

Before going on to study the behavior that can be expected to emerge, it is a good idea to take a closer look at which features of reality the Prisoner's Dilemma framework is, and is not, able to encompass. Fortunately, the very simplicity of the framework makes it possible to avoid many restrictive assumptions that would otherwise limit the analysis:

1. The payoffs of the players need not be comparable at all. For example, a journalist might get rewarded with another inside story, while the cooperating bureaucrat might be rewarded with a chance to have a policy argument presented in a favorable light.

2. The payoffs certainly do not have to be symmetric. It is a convenience to think of the interaction as exactly equivalent from the perspective of the two players, but this is not necessary. One does not have to assume, for example, that the reward for mutual cooperation, or any of the other three payoff parameters, have the same magnitude for both players. As mentioned earlier, one does not even have to assume that they are measured in comparable units. The only thing that has to be assumed is that, for each player, the four payoffs are ordered as required for the definition of the Prisoner's Dilemma.

3. The payoffs of a player do not have to be measured on an absolute scale. They need only be measured relative to each other.[6]

4. Cooperation need not be considered desirable from the point of view of the rest of the world. There are times when one wants to retard, rather than foster, cooperation between players. Collusive business practices are good for

the businesses involved but not so good for the rest of society. In fact, most forms of corruption are welcome instances of cooperation for the participants but are unwelcome to everyone else. So, on occasion, the theory will be used in reverse to show how to prevent, rather than to promote, cooperation.

5. There is no need to assume that the players are rational. They need not be trying to maximize their rewards. Their strategies may simply reflect standard operating procedures, rules of thumb, instincts, habits, or imitation (Simon 1955; Cyert and March 1963).

6. The actions that players take are not necessarily even conscious choices. A person who sometimes returns a favor, and sometimes does not, may not think about what strategy is being used. There is no need to assume deliberate choice at all.[7]

The framework is broad enough to encompass not only people but also nations and bacteria. Nations certainly take actions which can be interpreted as choices in a Prisoner's Dilemma—as in the raising or lowering of tariffs. It is not necessary to assume that such actions are rational or are the outcome of a unified actor pursuing a single goal. On the contrary, they might well be the result of an incredibly complex bureaucratic politics involving complicated information processing and shifting political coalitions (Allison 1971).

Likewise, at the other extreme, an organism does not need a brain to play a game. Bacteria, for example, are highly responsive to selected aspects of their chemical environment. They can therefore respond differentially to what other organisms are doing, and these conditional strategies of behavior can be inherited. Moreover, the behavior of a bacterium can affect the fitness of other organisms around

it, just as the behavior of other organisms can affect the fitness of a bacterium. But biological applications are best saved for chapter 5.

For now the main interest will be in people and organizations. Therefore, it is good to know that for the sake of generality, it is not necessary to assume very much about how deliberate and insightful people are. Nor is it necessary to assume, as the sociobiologists do, that important aspects of human behavior are guided by one's genes. The approach here is strategic rather than genetic.

Of course, the abstract formulation of the problem of cooperation as a Prisoner's Dilemma puts aside many vital features that make any actual interaction unique. Examples of what is left out by this formal abstraction include the possibility of verbal communication, the direct influence of third parties, the problems of implementing a choice, and the uncertainty about what the other player actually did on the preceding move. In chapter 8 some of these complicating factors are added to the basic model. It is clear that the list of potentially relevant factors that have been left out could be extended almost indefinitely. Certainly, no intelligent person should make an important choice without trying to take such complicating factors into account. The value of an analysis without them is that it can help to clarify some of the subtle features of the interaction—features which might otherwise be lost in the maze of complexities of the highly particular circumstances in which choice must actually be made. It is the very complexity of reality which makes the analysis of an abstract interaction so helpful as an aid to understanding.

The next chapter explores the emergence of cooperation through a study of what is a good strategy to employ if confronted with an iterated Prisoner's Dilemma. This exploration has been done in a novel way, with a computer tournament. Professional game theorists were invited to submit their favorite strategy, and each of these decision rules was paired off with each of the others to see which would do best overall. Amazingly enough, the winner was the simplest of all strategies submitted. This was TIT FOR TAT, the strategy which cooperates on the first move and then does whatever the other player did on the previous move. A second round of the tournament was conducted in which many more entries were submitted by amateurs and professionals alike, all of whom were aware of the results of the first round. The result was another victory for TIT FOR TAT! The analysis of the data from these tournaments reveals four properties which tend to make a decision rule successful: avoidance of unnecessary conflict by cooperating as long as the other player does, provocability in the face of an uncalled for defection by the other, forgiveness after responding to a provocation, and clarity of behavior so that the other player can adapt to your pattern of action.

These results from the tournaments demonstrate that under suitable conditions, cooperation can indeed emerge in a world of egoists without central authority. To see just how widely these results apply, a theoretical approach is taken in chapter 3. A series of propositions are proved that not only demonstrate the requirements for the emergence of cooperation but also provide the chronological story of the evolution of cooperation. Here is the argument in a nutshell. The evolution of cooperation requires that individuals have a sufficiently large chance to meet again so that they have a stake in their future interaction. If this is true, cooperation can evolve in three stages.

1. The beginning of the story is that cooperation can get started even in a world of unconditional defection. The development *cannot* take place if it is tried only by scattered individuals who have virtually no chance to interact with each other. However, cooperation can evolve from small clusters of individuals who base their cooperation on reciprocity and have even a small proportion of their interactions with each other.

2. The middle of the story is that a strategy based on reciprocity can thrive in a world where many different kinds of strategies are being tried.

3. The end of the story is that cooperation, once established on the basis of reciprocity, can protect itself from invasion by less cooperative strategies. Thus, the gear wheels of social evolution have a ratchet.

Chapters 4 and 5 take concrete settings to demonstrate just how widely these results apply. Chapter 4 is devoted to the fascinating case of the "live and let live" system which emerged during the trench warfare of World War I. In the midst of this bitter conflict, the front-line soldiers often refrained from shooting to kill—provided their restraint was reciprocated by the soldiers on the other side. What made this mutual restraint possible was the static nature of trench warfare, where the same small units faced each other for extended periods of time. The soldiers of these opposing small units actually violated orders from their own high commands in order to achieve tacit cooperation with each other. A detailed look at this case shows that when the conditions are present for the emergence of cooperation, cooperation can get started and prove stable in situations which otherwise appear extraordinarily unpromising. In particular, the "live and let live" system demonstrates that friendship is hardly necessary for the development of cooperation. Under suitable conditions, cooperation based upon reciprocity can develop even between antagonists.

Chapter 5, written with evolutionary biologist William D. Hamilton, demonstrates that cooperation can emerge even without foresight. This is done by showing that Cooperation Theory can account for the patterns of behavior found in a wide range of biological systems, from bacteria to birds. Cooperation in biological systems can occur even when the participants are not related, and even when they are unable to appreciate the consequences of their own behavior. What makes this possible are the evolutionary mechanisms of genetics and survival of the fittest. An individual able to achieve a beneficial response from another is more likely to have offspring that survive and that continue the pattern of behavior which elicited beneficial responses from others. Thus, under suitable conditions, cooperation based upon reciprocity proves stable in the biological world. Potential applications are spelled out for specific aspects of territoriality, mating, and disease. The conclusion is that Darwin's emphasis on individual advantage can, in fact, account for the presence of cooperation between individuals of the same or even different species. As long as the proper conditions are present, cooperation can get started, thrive, and prove stable.

While foresight is not necessary for the evolution of cooperation, it can certainly be helpful. Therefore chapters 6 and 7 are devoted to offering advice to participants and reformers, respectively. Chapter 6 spells out the implications of Cooperation Theory for anyone who is in a Prisoner's Dilemma. From the participant's point of view, the object is to do as well as possible, regardless of how well the other player does. Based upon the tournament results and the formal propositions, four simple suggestions are

offered for individual choice: do not be envious of the other player's success; do not be the first to defect; reciprocate both cooperation and defection; and do not be too clever.

Understanding the perspective of a participant can also serve as the foundation for seeing what can be done to make it easier for cooperation to develop among egoists. Thus, chapter 7 takes the Olympian perspective of a reformer who wants to alter the very terms of the interactions so as to promote the emergence of cooperation. A wide variety of methods are considered, such as making the interactions between the players more durable and frequent, teaching the participants to care about each other, and teaching them to understand the value of reciprocity. This reformer's perspective provides insights into a wide variety of topics, from the strength of bureaucracy to the difficulties of Gypsies, and from the morality of TIT FOR TAT to the art of writing treaties.

Chapter 8 extends the implications of Cooperation Theory into new domains. It shows how different kinds of social structure affect the way cooperation can develop. For example, people often relate to each other in ways that are influenced by observable features, such as sex, age, skin color, and style of dress. These cues can lead to social structures based on stereotyping and status hierarchies. As another example of social structure, the role of reputation is considered. The struggle to establish and maintain one's reputation can be a major feature of intense conflicts. For example, the American government's escalation of the war in Vietnam in 1965 was mainly due to its desire to deter other challenges to its interests by maintaining its reputation on the world stage. This chapter also considers a government's concern for maintaining its reputation with its own citizens. To be effective, a government cannot enforce any standards it chooses but must elicit compliance from a majority of the governed. To do this requires setting the rules so that most of the governed find it profitable to obey most of the time. The implications of this approach are fundamental to the operation of authority, and are illustrated by the regulation of industrial pollution and the supervision of divorce settlements.

By the final chapter, the discussion has developed from the study of the emergence of cooperation among egoists without central authority to an analysis of what happens when people actually *do* care about each other and what happens when there *is* central authority. But the basic approach is always the same: seeing how individuals operate in their own interest reveals what happens to the whole group. This approach allows more than the understanding of the perspective of a single player. It also provides an appreciation of what it takes to promote the stability of mutual cooperation in a given setting. The most promising finding is that if the facts of Cooperation Theory are known by participants with foresight, the evolution of cooperation can be speeded up.

II

The Emergence of Cooperation

CHAPTER 2

The Success of TIT FOR TAT in Computer Tournaments

SINCE the Prisoner's Dilemma is so common in everything from personal relations to international relations, it would be useful to know how best to act when in this type of setting. However, the proposition of the previous chapter demonstrates that there is no one best strategy to use. What is best depends in part on what the other player is likely to be doing. Further, what the other is likely to be doing may well depend on what the player expects *you* to do.

To get out of this tangle, help can be sought by combing the research already done concerning the Prisoner's Dilemma for useful advice. Fortunately, a great deal of research has been done in this area.

Psychologists using experimental subjects have found that, in the iterated Prisoner's Dilemma, the amount of cooperation attained—and the specific pattern for attaining it—depend on a wide variety of factors relating to the context of the game, the attributes of the individual players, and the relationship between the players. Since behavior in the game reflects so many important factors about people, it has become a standard way to explore questions in social psychology, from the effects of westernization in Central Africa (Bethlehem 1975) to the existence (or nonexistence) of aggression in career-oriented women (Baefsky and Berger 1974), and to the differential consequences of abstract versus concrete thinking styles (Nydegger 1974). In the last fifteen years, there have been hundreds of articles on the Prisoner's Dilemma cited in *Psychological Abstracts.* The iterated Prisoner's Dilemma has become the *E. coli* of social psychology.

Just as important as its use as an experimental test bed is the use of the Prisoner's Dilemma as the conceptual foundation for models of important social processes. Richardson's model of the arms race is based on an interaction which is essentially a Prisoner's Dilemma, played once a year with the budgets of the competing nations (Richardson 1960; Zinnes 1976, pp. 330–40). Oligopolistic competition can also be modeled as a Prisoner's Dilemma (Samuelson 1973, pp. 503–5). The ubiquitous problems of collective action to produce a collective good are analyzable as Prisoner's Dilemmas with many players (G. Hardin 1982). Even vote trading has been modeled as a Prisoner's Dilemma (Riker and Brams 1973). In fact, many of the best-developed models of important political, social, and economic processes have the Prisoner's Dilemma as their foundation.

There is yet a third literature about the Prisoner's Dilemma. This literature goes beyond the empirical questions of the laboratory or the real world, and instead uses the abstract game to analyze the features of some fundamental strategic issues, such as the meaning of rationality (Luce and Raiffa 1957), choices which affect other people (Schelling 1973), and cooperation without enforcement (Taylor 1976).

Unfortunately, none of these three literatures on the Prisoner's Dilemma reveals very much about how to play the game well. The experimental literature is not much help, because virtually all of it is based on analyzing the choices made by players who are seeing the formal game for the first time. Their appreciation of the strategic subtleties is bound to be restricted. Although the experimental subjects may have plenty of experience with everyday occurrences of the Prisoner's Dilemma, their ability to call on this experience in a formal setting may be limited. The choices of experienced economic and political elites in natural settings are studied in some of the applied literature of Prisoner's Dilemma, but the evidence is of limited help because of the relatively slow pace of most high-level interactions and the difficulty of controlling for changing circumstances. All together, no more than a few dozen choices have been identified and analyzed this way. Finally, the abstract literature of strategic interaction usually studies variants of the iterated Prisoner's Dilemma designed to eliminate the dilemma itself by introducing changes in the game, such as allowing interdependent choices (Howard 1966; Rapoport 1967), or putting a tax on defection (Tideman and Tullock 1976; Clarke 1980).

To learn more about how to choose effectively in an iterated Prisoner's Dilemma, a new approach is needed. Such an approach would have to draw on people who have a rich understanding of the strategic possibilities inherent in a non-zero-sum setting, a situation in which the interests of the participants partially coincide and partially conflict. Two important facts about non-zero-sum settings would have to be taken into account. First, the proposition of the previous chapter demonstrates that what is effective depends not only upon the characteristics of a particular strategy, but also upon the nature of the other strategies with which it must interact. The second point follows directly from the first. An effective strategy must be able at any point to take into account the history of the interaction as it has developed so far.

A computer tournament for the study of effective choice in the iterated Prisoner's Dilemma meets these needs. In a computer tournament, each entrant writes a program that embodies a rule to select the cooperative or noncooperative choice on each move. The program has available to it the history of the game so far, and may use this history in making a choice. If the participants are recruited primarily from those who are familiar with the Prisoner's Dilemma, the entrants can be assured that their decision rule will be facing rules of other informed entrants. Such recruitment would also guarantee that the state of the art is represented in the tournament.

Wanting to find out what would happen, I invited professional game theorists to send in entries to just such a computer tournament. It was structured as a round robin, meaning that each entry was paired with each other entry. As announced in the rules of the tournament, each entry was also paired with its own twin and with RANDOM, a program that randomly cooperates and defects with equal probability. Each game consisted of exactly two hundred moves.[1] The payoff matrix for each move was the familiar

one described in chapter 1. It awarded both players 3 points for mutual cooperation, and 1 point for mutual defection. If one player defected while the other player cooperated, the defecting player received 5 points and the cooperating player received 0 points.

No entry was disqualified for exceeding the allotted time. In fact, the entire round robin tournament was run five times to get a more stable estimate of the scores for each pair of players. In all, there were 120,000 moves, making for 240,000 separate choices.

The fourteen submitted entries came from five disciplines: psychology, economics, political science, mathematics, and sociology. Appendix A lists the names and affiliations of the people who submitted these entries, and it gives the rank and score of their entries.

One remarkable aspect of the tournament was that it allowed people from different disciplines to interact with each other in a common format and language. Most of the entrants were recruited from those who had published articles on game theory in general or the Prisoner's Dilemma in particular.

TIT FOR TAT, submitted by Professor Anatol Rapoport of the University of Toronto, won the tournament. This was the simplest of all submitted programs and it turned out to be the best!

TIT FOR TAT, of course, starts with a cooperative choice, and thereafter does what the other player did on the previous move. This decision rule is probably the most widely known and most discussed rule for playing the Prisoner's Dilemma. It is easily understood and easily programmed. It is known to elicit a good degree of cooperation when played with humans (Oskamp 1971; W. Wilson 1971). As an entry in a computer tournament, it has the desirable properties that it is not very exploitable and that it does well with its own twin. It has the disadvantage that it is too generous with the RANDOM rule, which was known by the participants to be entered in the tournament.

In addition, TIT FOR TAT was known to be a powerful competitor. In a preliminary tournament, TIT FOR TAT scored second place; and in a variant of that preliminary tournament, TIT FOR TAT won first place. All of these facts were known to most of the people designing programs for the Computer Prisoner's Dilemma Tournament, because they were sent copies of a description of the preliminary tournament. Not surprisingly, many of them used the TIT FOR TAT principle and tried to improve upon it.

The striking fact is that *none* of the more complex programs submitted was able to perform as well as the original, simple TIT FOR TAT.

This result contrasts with computer chess tournaments, where complexity is obviously needed. For example, in the Second World Computer Chess Championships, the least complex program came in last (Jennings 1978). It was submitted by Johann Joss of the Eidgenossishe Technische Hochschule of Zurich, Switzerland, who also submitted an entry to the Computer Prisoner's Dilemma Tournament. His entry to the Prisoner's Dilemma Tournament was a small modification of TIT FOR TAT. But his modification, like the others, just lowered the performance of the decision rule.

Analysis of the results showed that neither the discipline of the author, the brevity of the program—nor its *length*—accounts for a rule's relative success. What does?

Before answering this question, a remark on the interpretation of numerical scores is in order. In a game of 200 moves, a useful benchmark for very good performance is

600 points, which is equivalent to the score attained by a player when both sides always cooperate with each other. A useful benchmark for very poor performance is 200 points, which is equivalent to the score attained by a player when both sides never cooperate with each other. Most scores range between 200 and 600 points, although scores from 0 to 1000 points are possible. The winner, TIT FOR TAT, averaged 504 points per game.

Surprisingly, there is a single property which distinguishes the relatively high-scoring entries from the relatively low-scoring entries. This is the property of being *nice,* which is to say never being the first to defect. (For the sake of analyzing this tournament, the definition of a nice rule will be relaxed to include rules which will not be the first to defect before the last few moves, say before move 199.)

Each of the eight top-ranking entries (or rules) is nice. None of the other entries is. There is even a substantial gap in the score between the nice entries and the others. The nice entries received tournament averages between 472 and 504, while the best of the entries that were not nice received only 401 points. Thus, not being the first to defect, at least until virtually the end of the game, was a property which, all by itself, separated the more successful rules from the less successful rules in this Computer Prisoner's Dilemma Tournament.

Each of the nice rules got about 600 points with each of the other seven nice rules and with its own twin. This is because when two nice rules play, they are sure to cooperate with each other until virtually the end of the game. Actually the minor variations in end-game tactics did not account for much variation in the scores.

Since the nice rules all got within a few points of 600 with each other, the thing that distinguished the relative rankings among the nice rules was their scores with the rules which are not nice. This much is obvious. What is not obvious is that the relative ranking of the eight top rules was largely determined by just two of the other seven rules. These two rules are *kingmakers* because they do not do very well for themselves, but they largely determine the rankings among the top contenders.

The most important kingmaker was based on an "outcome maximization" principle originally developed as a possible interpretation of what human subjects do in the Prisoner's Dilemma laboratory experiments (Downing 1975). This rule, called DOWNING, is a particularly interesting rule in its own right. It is well worth studying as an example of a decision rule which is based upon a quite sophisticated idea. Unlike most of the others, its logic is not just a variant of TIT FOR TAT. Instead it is based on a deliberate attempt to understand the other player and then to make the choice that will yield the best long-term score based upon this understanding. The idea is that if the other player does not seem responsive to what DOWNING is doing, DOWNING will try to get away with whatever it can by defecting. On the other hand, if the other player does seem responsive, DOWNING will cooperate. To judge the other's responsiveness, DOWNING estimates the probability that the other player cooperates after it (DOWNING) cooperates, and also the probability that the other player cooperates after DOWNING defects. For each move, it updates its estimate of these two conditional probabilities and then selects the choice which will maximize its own long-term payoff under the assumption that it has correctly modeled the other player. If the two conditional probabilities have similar values, DOWNING determines that it pays to defect, since the other player seems to be doing the same thing whether DOWNING cooperates or not. Conversely, if the other player tends to cooperate after a cooperation but not after a defection by DOWNING, then the other player seems responsive, and DOWNING will calculate that the best thing to do with a responsive player is to cooperate. Under certain circumstances, DOWNING will even determine that the best strategy is to alternate cooperation and defection.

At the start of a game, DOWNING does not know the values of these conditional probabilities for the other players. It assumes that they are both .5, but gives no weight to this estimate when information actually does come in during the play of the game.

This is a fairly sophisticated decision rule, but its implementation does have one flaw. By initially assuming that the other player is unresponsive, DOWNING is doomed to defect on the first two moves. These first two defections led many other rules to punish DOWNING, so things usually got off to a bad start. But this is precisely why DOWNING served so well as a kingmaker. First-ranking TIT FOR TAT and second-ranking TIDEMAN AND CHIERUZZI both reacted in such a way that DOWNING learned to expect that defection does not pay but that cooperation does. All of the other nice rules went downhill with DOWNING.

The nice rules did well in the tournament largely because they did so well with each other, and because there were enough of them to raise substantially each other's average score. As long as the other player did not defect, each of the nice rules was certain to continue cooperating until virtually the end of the game. But what happened if there was a defection? Different rules responded quite differently, and their response was important in determining their overall success. A key concept in this regard is the forgiveness of a decision rule. *Forgiveness* of a rule can be informally described as its propensity to cooperate in the moves after the other player has defected.[2]

Of all the nice rules, the one that scored lowest was also the one that was least forgiving. This is FRIEDMAN, a totally unforgiving rule that employs permanent retaliation. It is never the first to defect, but once the other defects even once, FRIEDMAN defects from then on. In contrast, the winner, TIT FOR TAT, is unforgiving for one move, but thereafter is totally forgiving of that defection. After one punishment, it lets bygones be bygones.

One of the main reasons why the rules that are not nice did not do well in the tournament is that most of the rules in the tournament were not very forgiving. A concrete illustration will help. Consider the case of JOSS, a sneaky rule that tries to get away with an occasional defection. This decision rule is a variation of TIT FOR TAT. Like TIT FOR TAT, it always defects immediately after the other player defects. But instead of always cooperating after the other player cooperates, 10 percent of the time it defects after the other player cooperates. Thus it tries to sneak in an occasional exploitation of the other player.

This decision rule seems like a fairly small variation of TIT FOR TAT, but in fact its overall performance was much worse, and it is interesting to see exactly why. Table 1 shows the move-by-move history of a game between JOSS and TIT FOR TAT. At first both players cooperated, but on the sixth move, JOSS selected one of its probabilistic defections. On the next move JOSS cooperated again, but TIT FOR TAT defected in response to JOSS's previous defection. Then JOSS defected in response to TIT FOR

TABLE 1

Illustrative Game Between TIT FOR TAT and JOSS

moves	1–20	11111	23232	32323	23232
moves	21–40	32324	44444	44444	44444
moves	41–60	44444	44444	44444	44444
moves	61–80	44444	44444	44444	44444
moves	81–100	44444	44444	44444	44444
moves	101–120	44444	44444	44444	44444
moves	121–140	44444	44444	44444	44444
moves	141–160	44444	44444	44444	44444
moves	161–180	44444	44444	44444	44444
moves	181–200	44444	44444	44444	44444

Score in this game: TIT FOR TAT 236; JOSS 241.
Legend: 1 both cooperated
2 TIT FOR TAT only cooperated
3 JOSS only cooperated
4 neither cooperated

TAT's defection. In effect, the single defection of JOSS on the sixth move created an *echo* back and forth between JOSS and TIT FOR TAT. This echo resulted in JOSS defecting on all the subsequent even numbered moves and TIT FOR TAT defecting on all the subsequent odd numbered moves.

On the twenty-fifth move, JOSS selected another of its probabilistic defections. Of course, TIT FOR TAT defected on the very next move and another reverberating echo began. This echo had JOSS defecting on the odd numbered moves. Together these two echoes resulted in both players defecting on every move after move 25. This string of mutual defections meant that for the rest of the game they both got only one point per turn. The final score of this game was 236 for TIT FOR TAT and 241 for JOSS. Notice that while JOSS did a little better than TIT FOR TAT, both did poorly.[3]

The problem was a combination of an occasional defection after the other's cooperation by JOSS, combined with a short-term lack of forgiveness by both sides. The moral is that if both sides retaliate in the way that JOSS and TIT FOR TAT did, it does not pay to be as greedy as JOSS was.

A major lesson of this tournament is the importance of minimizing echo effects in an environment of mutual power. When a single defection can set off a long string of recriminations and counterrecriminations, both sides suffer. A sophisticated analysis of choice must go at least three levels deep to take account of these echo effects. The first level of analysis is the direct effect of a choice. This is easy, since a defection always earns more than a cooperation. The second level considers the indirect effects, taking into account that the other side may or may not punish a defection. This much of the analysis was certainly appreciated by many of the entrants. But the third level goes deeper and takes into account the fact that in responding to the defections of the other side, one may be repeating or even amplifying one's own previous exploitative choice. Thus a single defection may be successful when analyzed for its direct effects, and perhaps even when its secondary effects are taken into account. But the real costs may be in the tertiary effects when one's own isolated defections turn into unending mutual recriminations. Without their realizing it, many of these rules actually wound up punishing themselves. With the other player serving as a mechanism to delay the self-punishment by a few moves, this aspect of self-punishment was not picked up by many of the decision rules.

Despite the fact that none of the attempts at more or less sophisticated decision rules was an improvement on TIT FOR TAT, it was easy to find several rules that would have performed substantially better than TIT FOR TAT in the

environment of the tournament. The existence of these rules should serve as a warning against the facile belief that an eye for an eye is necessarily the best strategy. There are at least three rules that would have won the tournament if submitted.

The sample program sent to prospective contestants to show them how to make a submission would in fact have won the tournament if anyone had simply clipped it and mailed it in! But no one did. The sample program defects only if the other player defected on the previous two moves. It is a more forgiving version of TIT FOR TAT in that it does not punish isolated defections. The excellent performance of this TIT FOR TWO TATS rule highlights the fact that a common error of the contestants was to expect that gains could be made from being relatively less forgiving than TIT FOR TAT, whereas in fact there were big gains to be made from being even more forgiving. The implication of this finding is striking, since it suggests that even expert strategists do not give sufficient weight to the importance of forgiveness.

Another rule which would have won the tournament was also available to most of the contestants. This was the rule which won the preliminary tournament, a report of which was used in recruiting the contestants. Called LOOK AHEAD, it was inspired by techniques used in artificial intelligence programs to play chess. It is interesting that artificial intelligence techniques could have inspired a rule which was in fact better than any of the rules designed by game theorists specifically for the Prisoner's Dilemma.

A third rule which would have won the tournament was a slight modification of DOWNING. If DOWNING had started with initial assumptions that the other players would be responsive rather than unresponsive, it too would have won and won by a large margin. A kingmaker could have been king. DOWNING's initial assumptions about the other players were pessimistic. It turned out that optimism about their responsiveness would not only have been more accurate but would also have led to more successful performance. It would have resulted in first place rather than tenth place.[4]

These results from supplementary rules reinforce a theme from the analysis of the tournament entries themselves: the entries were too competitive for their own good. In the first place, many of them defected early in the game without provocation, a characteristic which was very costly in the long run. In the second place, the optimal amount of forgiveness was considerably greater than displayed by any of the entries (except possibly DOWNING). And in the third place, the entry that was most different from the others, DOWNING, floundered on its own misplaced pessimism regarding the initial responsiveness of the others.

The analysis of the tournament results indicate that there is a lot to be learned about coping in an environment of mutual power. Even expert strategists from political science, sociology, economics, psychology, and mathematics made the systematic errors of being too competitive for their own good, not being forgiving enough, and being too pessimistic about the responsiveness of the other side.

The effectiveness of a particular strategy depends not only on its own characteristics, but also on the nature of the other strategies with which it must interact. For this reason, the results of a single tournament are not definitive. Therefore, a second round of the tournament was conducted.

The results of the second round provide substantially

better grounds for insight into the nature of effective choice in the Prisoner's Dilemma. The reason is that the entrants to the second round were all given the detailed analysis of the first round, including a discussion of the supplemental rules that would have done very well in the environment of the first round. Thus they were aware not only of the outcome of the first round, but also of the concepts used to analyze success, and the strategic pitfalls that were discovered. Moreover, they each knew that the others knew these things. Therefore, the second round presumably began at a much higher level of sophistication than the first round, and its results could be expected to be that much more valuable as a guide to effective choice in the Prisoner's Dilemma.

The second round was also a dramatic improvement over the first round in sheer size of the tournament. The response was far greater than anticipated. There was a total of sixty-two entries from six countries. The contestants were largely recruited through announcements in journals for users of small computers. The game theorists who participated in the first round of the tournament were also invited to try again. The contestants ranged from a ten-year-old computer hobbyist to professors of computer science, physics, economics, psychology, mathematics, sociology, political science, and evolutionary biology. The countries represented were the United States, Canada, Great Britain, Norway, Switzerland, and New Zealand.

The second round provided a chance both to test the validity of the themes developed in the analysis of the first round and to develop new concepts to explain successes and failures. The entrants also drew their own lessons from the experience of the first round. But different people drew different lessons. What is particularly illuminating in the second round is the way the entries based on different lessons actually interact.

TIT FOR TAT was the simplest program submitted in the first round, and it won the first round. It was the simplest submission in the second round, and it won the second round. Even though all the entrants to the second round knew that TIT FOR TAT had won the first round, no one was able to design an entry that did any better.

This decision rule was known to all of the entrants to the second round because they all had the report of the earlier round, showing that TIT FOR TAT was the most successful rule so far. They had read the arguments about how it was known to elicit a good degree of cooperation when played with humans, how it is not very exploitable, how it did well in the preliminary tournament, and how it won the first round. The report on the first round also explained some of the reasons for its success, pointing in particular to its property of never being the first to defect ("niceness") and its propensity to cooperate after the other player defected ("forgiveness" with the exception of a single punishment).

Even though an explicit tournament rule allowed anyone to submit any program, even one authored by someone else, only one person submitted TIT FOR TAT. This was Anatol Rapoport, who submitted it the first time.

The second round of the tournament was conducted in the same manner as the first round, except that minor end-game effects were eliminated. As announced in the rules, the length of the games was determined probabilistically with a 0.00346 chance of ending with each given move.[5] This is equivalent to setting $w = .99654$. Since no one knew

exactly when the last move would come, end-game effects were successfully avoided in the second round.

Once again, none of the personal attributes of the contestants correlated significantly with the performance of the rules. The professors did not do significantly better than the others, nor did the Americans. Those who wrote in FORTRAN rather than BASIC did not do significantly better either, even though the use of FORTRAN would usually indicate access to something more than a bottom-of-the-line microcomputer. The names of the contestants are shown in the order of their success in appendix A along with some information about them and their programs.

On average, short programs did not do significantly better than long programs, despite the victory of TIT FOR TAT. But on the other hand, neither did long programs (with their greater complexity) do any better than short programs.

The determination of what does account for success in the second round is not easy because there were 3969 ways the 63 rules (including RANDOM) were paired in the round robin tournament. This very large tournament score matrix is given in Appendix A along with information about the entrants and their programs. In all, there were over a million moves in the second round.

As in the first round, it paid to be nice. Being the first to defect was usually quite costly. More than half of the entries were nice, so obviously most of the contestants got the message from the first round that it did not pay to be the first to defect.

In the second round, there was once again a substantial correlation between whether a rule was nice and how well it did. Of the top fifteen rules, all but one were nice (and that one ranked eighth). Of the bottom fifteen rules, all but one were not nice. The overall correlation between whether a rule was nice and its tournament score was a substantial .58.

A property that distinguishes well among the nice rules themselves is how promptly and how reliably they responded to a challenge by the other player. A rule can be called *retaliatory* if it immediately defects after an "uncalled for" defection from the other. Exactly what is meant by "uncalled for" is not precisely determined. The point, however, is that unless a strategy is incited to an immediate response by a challenge from the other player, the other player may simply take more and more frequent advantage of such an easygoing strategy.

There were a number of rules in the second round of the tournament that deliberately used controlled numbers of defections to see what they could get away with. To a large extent, what determined the actual rankings of the nice rules was how well they were able to cope with these challengers. The two challengers that were especially important in this regard I shall called TESTER and TRANQUILIZER.

TESTER was submitted by David Gladstein and came in forty-sixth in the tournament. It is designed to look for softies, but is prepared to back off if the other player shows it won't be exploited. The rule is unusual in that it defects on the very first move in order to test the other's response. If the other player ever defects, it apologizes by cooperating and playing tit-for-tat for the rest of the game. Otherwise, it cooperates on the second and third moves but defects every other move after that. TESTER did a good job of exploiting several supplementary rules that would have

done quite well in the environment of the first round of the tournament. For example, TIT FOR TWO TATS defects only after the other player defects on the preceding two moves. But TESTER never does defect twice in a row. So TIT FOR TWO TATS always cooperates with TESTER, and gets badly exploited for its generosity. Notice that TESTER itself did not do particularly well in the tournament. It did, however, provide low scores for some of the more easygoing rules.

As another example of how TESTER causes problems for some rules which had done well in the first round, consider the three variants of Leslie Downing's outcome maximization principle. There were two separate submissions of the REVISED DOWNING program, based on DOWNING, which looked so promising in round one. These came from Stanley F. Quayle and Leslie Downing himself. A slightly modified version came from a youthful competitor, eleven-year-old Steve Newman. However, all three were exploited by TESTER since they all calculated that the best thing to do with a program that cooperated just over half the time after one's own cooperation was to keep on cooperating. Actually they would have been better off doing what TIT FOR TAT and many other high-ranking programs did, which was to defect immediately on the second move in response to TESTER's defection on the first move. This would have elicited TESTER's apology and things would have gone better thereafter.

TRANQUILIZER illustrates a more subtle way of taking advantage of many rules, and hence a more subtle challenge. It first seeks to establish a mutually rewarding relationship with the other player, and only then does it cautiously try to see if it will be allowed to get away with something. TRANQUILIZER was submitted by Craig

Feathers and came in twenty-seventh in the tournament. The rule normally cooperates but is ready to defect if the other player defects too often. Thus the rule tends to cooperate for the first dozen or two dozen moves if the other player is cooperating. Only then does it throw in an unprovoked defection. By waiting until a pattern of mutual cooperation has been developed, it hopes to lull the other side into being forgiving of occasional defections. If the other player continues to cooperate, then defections become more frequent. But as long as TRANQUILIZER is maintaining an average payoff of at least 2.25 points per move, it does not defect twice in succession, and it does not defect more than one-quarter of the time. It tries to avoid pressing its luck too far.

What it takes to do well with challenging rules like TESTER and TRANQUILIZER is to be ready to retaliate after an "uncalled for" defection from the other. So while it pays to be nice, it also pays to be retaliatory. TIT FOR TAT combines these desirable properties. It is nice, forgiving, and retaliatory. It is never the first to defect; it forgives an isolated defection after a single response; but it is always incited by a defection no matter how good the interaction has been so far.

The lessons of the first round of the tournament affected the environment of the second round, since the contestants were familiar with the results. The report on the first round of the Computer Prisoner's Dilemma Tournament (Axelrod 1980*a*) concluded that it paid to be not only nice but also forgiving. The contestants in the second round knew that such forgiving decision rules as TIT FOR TWO TATS and REVISED DOWNING would have done even better than TIT FOR TAT in the environment of the first round.

In the second round, many contestants apparently hoped that these conclusions would still be relevant. Of the sixty-two entries, thirty-nine were nice, and nearly all of them were at least somewhat forgiving. TIT FOR TWO TATS itself was submitted by an evolutionary biologist from the United Kingdom, John Maynard Smith. But it came in only twenty-fourth. As mentioned earlier, REVISED DOWNING was submitted twice. But in the second round, it was in the bottom half of the tournament.

What seems to have happened is an interesting interaction between people who drew one lesson and people who drew another from the first round. Lesson One was: "Be nice and forgiving." Lesson Two was more exploitative: "If others are going to be nice and forgiving, it pays to try to take advantage of them." The people who drew Lesson One suffered in the second round from those who drew Lesson Two. Rules like TRANQUILIZER and TESTER were effective at exploiting rules which were too easygoing. But the people who drew Lesson Two did not themselves do very well either. The reason is that in trying to exploit other rules, they often eventually got punished enough to make the whole game less rewarding for *both* players than pure mutual cooperation would have been. For example, TRANQUILIZER and TESTER themselves achieved only twenty-seventh and forty-sixth place, respectively. Each surpassed TIT FOR TAT's score with fewer than one-third of the rules. None of the other entries that tried to apply the exploitative conclusion of Lesson Two ranked near the top either.

While the use of Lesson Two tended to invalidate Lesson One, no entrants were able to benefit more than they were hurt in the tournament by their attempt to exploit the easygoing rules. The most successful entries tended to be relatively small variations on TIT FOR TAT which were designed to recognize and give up on a seemingly RANDOM player or a very uncooperative player. But the implementations of these ideas did not do better than the pure form of TIT FOR TAT. So TIT FOR TAT, which got along with almost everyone, won the second round of the tournament just as it had won the first round.

Would the results of the second round have been much different if the distribution of entries had been substantially different? Put another way, does TIT FOR TAT do well in a wide variety of environments? That is to say, is it *robust?*

A good way to examine this question is to construct a series of hypothetical tournaments, each with a very different distribution of the types of rules participating. The method of constructing these drastically modified tournaments is explained in appendix A. The results were that TIT FOR TAT won five of the six major variants of the tournament, and came in second in the sixth. This is a strong test of how robust the success of TIT FOR TAT really is.

Another way to examine the robustness of the results is to construct a whole sequence of hypothetical future rounds of the tournament. Some of the rules were so unsuccessful that they would be unlikely to be tried again in future tournaments, while others were successful enough that their continued presence in later tournaments would be likely. For this reason, it would be helpful to analyze what would happen over a series of tournaments if the more successful rules became a larger part of the environment for each rule, and the less successful rules were met less often. This analysis would be a strong test of a rule's performance, because continued success would require a rule to do well with other successful rules.

Evolutionary biology provides a useful way to think about this dynamic problem (Trivers 1971; Dawkins 1976, pp. 197–202; Maynard Smith 1978). Imagine that there are many animals of a single species which interact with each other quite often. Suppose the interactions take the form of a Prisoner's Dilemma. When two animals meet, they can cooperate with each other, not cooperate with each other, or one animal could exploit the other. Suppose further that each animal can recognize individuals it has already interacted with and can remember salient aspects of their interaction, such as whether the other has usually cooperated. A round of the tournament can then be regarded as a simulation of a single generation of such animals, with each decision rule being employed by large numbers of individuals. One convenient implication of this interpretation is that a given animal can interact with another animal using its own decision rule, just as it can run into an animal using some other rule.

The value of this analogy is that it allows a simulation of future generations of a tournament. The idea is that the more successful entries are more likely to be submitted in the next round, and the less successful entries are less likely to be submitted again. To make this precise, we can say that the number of copies (or offspring) of a given entry will be proportional to that entry's tournament score. We simply have to interpret the average payoff received by an individual as proportional to the individual's expected number of offspring. For example, if one rule gets twice as high a tournament score in the initial round as another rule, then it will be twice as well-represented in the next round.[6] Thus, RANDOM, for example, will be less important in the second generation, whereas TIT FOR TAT and the other high-ranking rules will be better represented.

In human terms, a rule which was not scoring well might be less likely to appear in the future for several different reasons. One possibility is that a player will try different strategies over time, and then stick with what seems to work best. Another possibility is that a person using a rule sees that other strategies are more successful and therefore switches to one of those strategies. Still another possibility is that a person occupying a key role, such as a member of Congress or the manager of a business, would be removed from that role if the strategy being followed was not very successful. Thus, learning, imitation, and selection can all operate in human affairs to produce a process which makes relatively unsuccessful strategies less likely to appear later.

The simulation of this process for the Prisoner's Dilemma tournament is actually quite straightforward. The tournament matrix gives the score each strategy gets with each of the other strategies. Starting with the proportions of each type in a given generation, it is only necessary to calculate the proportions which will exist in the next generation.[7] The better a strategy does, the more its representation will grow.

The results provide an interesting story. The first thing that happens is that the lowest-ranking eleven entries fall to half their initial size by the fifth generation while the middle-ranking entries tend to hold their own and the top-ranking entries slowly grow in size. By the fiftieth generation, the rules that ranked in the bottom third of the tournament have virtually disappeared, while most of those in the middle third have started to shrink, and those in the top third are continuing to grow (see figure 2).

This process simulates survival of the fittest. A rule that is successful on average with the current distribution of

FIGURE 2

Simulated Ecological Success of the Decision Rules

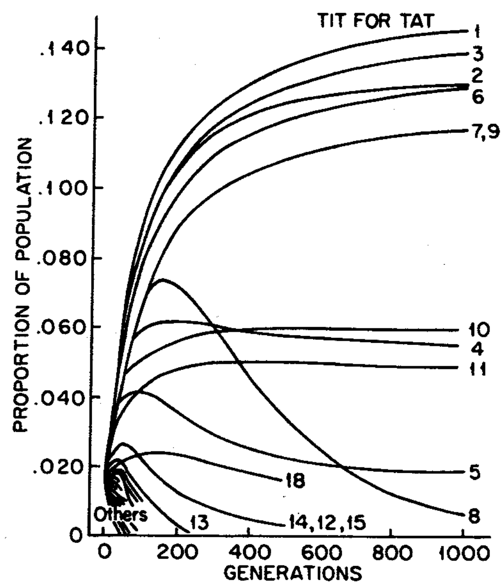

rules in the population will become an even larger proportion of the environment of the other rules in the next generation. At first, a rule that is successful with all sorts of rules will proliferate, but later as the unsuccessful rules disappear, success requires good performance with other successful rules.

This simulation provides an ecological perspective because there are no new rules of behavior introduced. It differs from an evolutionary perspective, which would allow mutations to introduce new strategies into the environment. In the ecological perspective there is a changing distribution of given types of rules. The less successful rules become less common and the more successful rules proliferate. The statistical distribution of types of individuals changes in each generation, and this changes the environment with which each of the individual types has to interact.

At first, poor programs and good programs are represented in equal proportions. But as time passes, the poorer ones begin to drop out and the good ones thrive. Success breeds more success, provided that the success derives from interactions with other successful rules. If, on the other hand, a decision rule's success derives from its ability to exploit other rules, then as these exploited rules die out, the exploiter's base of support becomes eroded and the exploiter suffers a similar fate.

A good example of ecological extinction is provided by HARRINGTON, the only non-nice rule among the top fifteen finishers in the second round. In the first two hundred or so generations of the ecological tournament, as TIT FOR TAT and the other successful nice programs were increasing their percentage of the population, HARRINGTON was also increasing its percentage. This was because of HARRINGTON's exploitative strategy. By the two hundredth generation or so, things began to take a noticeable turn. Less successful programs were becoming extinct, which meant that there were fewer and fewer prey for HARRINGTON to exploit. Soon HARRINGTON could not keep up with the successful nice rules, and by the one thousandth generation HARRINGTON was as extinct as the exploitable rules on which it preyed.

The ecological analysis shows that doing well with rules that do not score well themselves is eventually a self-defeating process. Not being nice may look promising at first, but in the long run it can destroy the very environment it needs for its own success.

The results also provide yet another victory for TIT FOR TAT. TIT FOR TAT had a very slight lead in the

original tournament, and never lost this lead in simulated generations. By the one-thousandth generation it was the most successful rule and still growing at a faster rate than any other rule.

The overall record of TIT FOR TAT is very impressive. To recapitulate, in the second round, TIT FOR TAT achieved the highest average score of the sixty-two entries in the tournament. It also achieved the highest score in five of the six hypothetical tournaments which were constructed by magnifying the effects of different types of rules from the second round. And in the sixth hypothetical tournament it came in second. Finally, TIT FOR TAT never lost its first-place standing in a simulation of future generations of the tournament. Added to its victory in the first round of the tournament, and its fairly good performance in laboratory experiments with human subjects, TIT FOR TAT is clearly a very successful strategy.

Proposition 1 says that there is no absolutely best rule independent of the environment. What can be said for the empirical successes of TIT FOR TAT is that it is a very robust rule: it does very well over a wide range of environments. Part of its success might be that other rules anticipate its presence and are designed to do well with it. Doing well with TIT FOR TAT requires cooperating with it, and this in turn helps TIT FOR TAT. Even rules like TESTER that were designed to see what they could get away with, quickly apologize to TIT FOR TAT. Any rule which tries to take advantage of TIT FOR TAT will simply hurt itself. TIT FOR TAT benefits from its own nonexploitability because three conditions are satisfied:

1. The possibility of encountering TIT FOR TAT is salient.
2. Once encountered, TIT FOR TAT is easy to recognize.

3. Once recognized, TIT FOR TAT's nonexploitability is easy to appreciate.

Thus TIT FOR TAT benefits from its own *clarity.*

On the other hand, TIT FOR TAT foregoes the possibility of exploiting other rules. While such exploitation is occasionally fruitful, over a wide range of environments the problems with trying to exploit others are manifold. In the first place, if a rule defects to see what it can get away with, it risks retaliation from the rules that are provocable. In the second place, once mutual recriminations set in, it can be difficult to extract oneself. And, finally, the attempt to identify and give up on unresponsive rules (such as RANDOM or excessively uncooperative rules) often mistakenly led to giving up on rules which were in fact salvageable by a more patient rule like TIT FOR TAT. Being able to exploit the exploitable without paying too high a cost with the others is a task which was not successfully accomplished by any of the entries in round two of the tournament.

What accounts for TIT FOR TAT's robust success is its combination of being nice, retaliatory, forgiving, and clear. Its niceness prevents it from getting into unnecessary trouble. Its retaliation discourages the other side from persisting whenever defection is tried. Its forgiveness helps restore mutual cooperation. And its clarity makes it intelligible to the other player, thereby eliciting long-term cooperation.

CHAPTER 9

The Robustness of Reciprocity

THE EVOLUTIONARY APPROACH is based on a simple principle: whatever is successful is likely to appear more often in the future. The mechanism can vary. In classical Darwinian evolution, the mechanism is natural selection based upon differential survival and reproduction. In Congress, the mechanism can be an increased chance of reelection for those members who are effective in delivering legislation and services for their constituency. In the business world, the mechanism can be the avoidance of bankruptcy by a profitable company. But the evolutionary mechanism need not be a question of life and death. With intelligent players, a successful strategy can appear more often in the future because other players convert to it. The conversion can be based on more or less blind imitation of the successful players, or it can be based on a more or less informed process of learning.

The evolutionary process needs more than differential growth of the successful. In order to go very far it also needs a source of variety—of new things being tried. In the genetics of biology, this variety is provided by mutation and by a reshuffling of genes with each generation. In social processes, the variety can be introduced by the "trial" in "trial and error" learning. This kind of learning might or might not reflect a high degree of intelligence. A new pattern of behavior might be undertaken simply as a random variant of an old pattern of behavior, or the new strategy could be deliberately constructed on the basis of prior experience and a theory about what is likely to work best in the future.

To study different aspects of the evolutionary process, different methodological tools have been used. One set of questions asked about the destination of the evolutionary process. To study this, the concept of collective (or evolutionary) stability was used to study where the evolutionary process would stop. The idea was to determine which strategies could not be invaded if they were used by everyone. The virtue of this approach is that it allowed a good specification of which types of strategies can protect themselves, and under what conditions this protection can work. For example, it was shown that TIT FOR TAT would be collectively stable if the shadow of the future were large enough, and that the strategy of always defecting would be collectively stable under all possible conditions.

The power of the collective stability approach is that it allows a consideration of *all* possible new strategies, whether minor variants of the common strategy or completely new ideas. The limitation of the stability approach is that it only

tells what will last once established, but it does not tell what will get established in the first place. Since many different strategies can be collectively stable once established in a population, it is important to know which strategies are likely to get established in the first place. For this a different methodology was needed.

To see what is likely to get established in the first place, the emphasis must be placed upon the variety of things that can happen at once in a population. To capture this variety, the tournament approach was used. The tournament itself was conducted to encourage the presence of sophisticated strategies, which were attained in the first round by soliciting entries from professional game theorists. Refinement of strategies was carried further in the second round by making sure that the new entrants were aware of the results of the first round. Thus, new ideas could enter the tournament either as refinements of the old ideas or as totally new conceptions of what might work best. Then the analysis of what actually worked best in this variegated environment told a great deal about which kind of strategy is likely to flourish.

Since the process of getting fully established is likely to take a considerable amount of time, another kind of technique was used to study the changing prospects of strategies as their social environment changes. This technique was an ecological analysis, which calculated what would happen if each generation had strategies growing in frequency in proportion to their success in the previous generation. This was an ecological approach because it introduced no new strategies, but instead determined the consequences over hundreds of generations of the variety of strategies already represented in the tournament. It allowed for an analysis of whether the strategies that were successful in the beginning would remain successful after the poor performers had dropped out. The growth of the successful strategies in each generation could be thought of as due to either better survival and reproduction of the users of that strategy, or due to a greater chance of being imitated by the others.

Related to the ecological analysis was the territorial analysis of what would happen if the sixty-three strategies of the second round of the tournament were scattered in a territorial structure, with the player at each location interacting with the four neighbors of that location. In the territorial system, determination of what is successful is local. Each location which has a more successful neighbor adopts the strategy of the most successful of its neighbors. As in the ecological simulation, this growth of the more successful can be attributed to either better survival and reproduction, or to a greater chance of being imitated by others.

To use these tools of evolutionary analysis, what is needed is a way to determine how any given strategy will perform with any other given strategy. In simple cases, this calculation can be done algebraically, as in the determination of how TIT FOR TAT will do when it meets a player who always defects. In more complex cases, the calculation can be done by simulating the interactions and cumulating the payoffs received, as in the conduct of the Computer Tournament for the Prisoner's Dilemma. The ideas of a time discount and uncertain ending of the interaction were incorporated in the tournament by varying the lengths of the games. The consequences of the probabilistic nature of some strategies were handled by averaging over several repetitions of the interaction between the same two strategies.

These tools of evolutionary analysis could be used with any social setting. In this book they have been applied to

one particular kind of social setting, a setting which captures the fundamental dilemma of cooperation. The potential for cooperation arises when each player can help the other. The dilemma arises when giving this help is costly. The opportunity for mutual gain from cooperation comes into play when the gains from the other's cooperation are larger than the costs of one's own cooperation. In that case mutual cooperation is preferred by both to mutual noncooperation (so-called defection). But getting what you prefer is not so easy. There are two reasons. In the first place, you have to get the other player to help—even though the other player is better off in the short run by not helping. In the second place, you are tempted to get whatever help you can without providing any costly help yourself.[1]

The main results of Cooperation Theory are encouraging. They show that cooperation can get started by even a small cluster of individuals who are prepared to reciprocate cooperation, even in a world where no one else will cooperate. The analysis also shows that the two key requisites for cooperation to thrive are that the cooperation be based on reciprocity, and that the shadow of the future is important enough to make this reciprocity stable. But once cooperation based on reciprocity is established in a population, it can protect itself from invasion by uncooperative strategies.

It is encouraging to see that cooperation can get started, can thrive in a variegated environment, and can protect itself once established. But what is most interesting is how little had to be assumed about the individuals or the social setting to establish these results. The individuals do not have to be rational: the evolutionary process allows the successful strategies to thrive, even if the players do not know why or how. Nor do the players have to exchange messages or commitments: they do not need words, because their deeds speak for them. Likewise, there is no need to assume trust between the players: the use of reciprocity can be enough to make defection unproductive. Altruism is not needed: successful strategies can elicit cooperation even from an egoist. Finally, no central authority is needed: cooperation based on reciprocity can be self-policing.

The emergence, growth, and maintenance of cooperation do require some assumptions about the individuals and the social setting. They require an individual to be able to recognize another player who has been dealt with before. They also require that one's prior history of interactions with this player can be remembered, so that a player can be responsive. Actually, these requirements for recognition and recall are not as strong as they might seem. Even bacteria can fulfill them by interacting with only one other organism and using a strategy (such as TIT FOR TAT) which responds only to the recent behavior of the other player. And if bacteria can play games, so can people and nations.

For cooperation to prove stable, the future must have a sufficiently large shadow. This means that the importance of the next encounter between the same two individuals must be great enough to make defection an unprofitable strategy when the other player is provocable. It requires that the players have a large enough chance of meeting again and that they do not discount the significance of their next meeting too greatly. For example, what made cooperation possible in the trench warfare of World War I was the fact that the same small units from opposite sides of no-man's-land would be in contact for long periods of time, so that if one side broke the tacit understandings, then the other side could retaliate against the same unit.

Finally, the evolution of cooperation requires that suc-

cessful strategies can thrive and that there be a source of variation in the strategies which are being used. These mechanisms can be classical Darwinian survival of the fittest and the mutation, but they can also involve more deliberate processes such as imitation of successful patterns of behavior and intelligently designed new strategic ideas.

In order for cooperation to get started in the first place, one more condition is required.The problem is that in a world of unconditional defection, a single individual who offers cooperation cannot prosper unless others are around who will reciprocate. On the other hand, cooperation can emerge from small clusters of discriminating individuals as long as these individuals have even a small proportion of their interactions with each other. So there must be some clustering of individuals who use strategies with two properties: the strategies will be the first to cooperate, and they will discriminate between those who respond to the cooperation and those who do not.

The conditions for the evolution of cooperation tell what is necessary, but do not, by themselves, tell what strategies will be most successful. For this question, the tournament approach has offered striking evidence in favor of the robust success of the simplest of all discriminating strategies: TIT FOR TAT. By cooperating on the first move, and then doing whatever the other player did on the previous move, TIT FOR TAT managed to do well with a wide variety of more or less sophisticated decision rules. It not only won the first round of the Computer Prisoner's Dilemma Tournament when facing entries submitted by professional game theorists, but it also won the second round which included over sixty entries designed by people who were able to take the results of the first round into account. It was also the winner in five of the six major variants of the second round (and second in the sixth variant). And most impressive, its success was not based only upon its ability to do well with strategies which scored poorly for themselves. This was shown by an ecological analysis of hypothetical future rounds of the tournament. In this simulation of hundreds of rounds of the tournament, TIT FOR TAT again was the most successful rule, indicating that it can do well with good and bad rules alike.

TIT FOR TAT's robust success is due to being nice, provocable, forgiving, and clear. Its niceness means that it is never the first to defect, and this property prevents it from getting into unnecessary trouble. Its retaliation discourages the other side from persisting whenever defection is tried. Its forgiveness helps restore mutual cooperation. And its clarity makes its behavioral pattern easy to recognize; and once recognized, it is easy to perceive that the best way of dealing with TIT FOR TAT is to cooperate with it.

Despite its robust success, TIT FOR TAT cannot be called the ideal strategy to play in the iterated Prisoner's Dilemma. For one thing, TIT FOR TAT and other nice rules require for their effectiveness that the shadow of the future be sufficiently great. But even then there is no ideal strategy independent of the strategies used by the others. In some extreme environments, even TIT FOR TAT would do poorly—as would be the case if there were not enough others who would ever reciprocate its initial cooperative choice. And TIT FOR TAT does have its strategic weaknesses as well. For example, if the other player defects once, TIT FOR TAT will always respond with a defection, and then if the other player does the same in response, the result would be an unending echo of alternating defections. In this sense, TIT FOR TAT is not forgiving enough. But

another problem is that TIT FOR TAT is too forgiving to those rules which are totally unresponsive, such as a completely random rule. What can be said for TIT FOR TAT is that it does indeed perform well in a wide variety of settings where the other players are all using more or less sophisticated strategies which themselves are designed to do well.

If a nice strategy, such as TIT FOR TAT, does eventually come to be adopted by virtually everyone, then individuals using this nice strategy can afford to be generous in dealing with any others. In fact, a population of nice rules can also protect itself from clusters of individuals using any other strategy just as well as they can protect themselves against single individuals.

These results give a chronological picture for the evolution of cooperation. Cooperation can begin with small clusters. It can thrive with rules that are nice, provocable, and somewhat forgiving. And once established in a population, individuals using such discriminating strategies can protect themselves from invasion. The overall level of cooperation tends to go up and not down. In other words, the machinery for the evolution of cooperation contains a ratchet.

The operation of this ratchet was seen in the development of the norm of reciprocity in the United States Congress. As described in the first chapter, in the early days of the republic, members of Congress were known for their deceit and treachery. They were quite unscrupulous and frequently lied to each other. Yet, over the years, cooperative patterns of behavior emerged and proved stable. These patterns were based upon the norm of reciprocity.

Many other institutions have developed stable patterns of cooperation based upon similar norms. Diamond markets, for example, are famous for the way their members exchange millions of dollars worth of goods with only a verbal pledge and a handshake. The key factor is that the participants know they will be dealing with each other again and again. Therefore any attempt to exploit the situation will simply not pay.

A wonderful illustration of this principle is provided in the memoirs of Ron Luciano, a baseball umpire who sometimes had his "bad days."

> Over a period of time I learned to trust certain catchers so much that I actually let them umpire for me on the bad days. The bad days usually followed the good nights. . . . On those days there wasn't much I could do but take two aspirins and call as little as possible. If someone I trusted was catching . . . I'd tell them, "Look, it's a bad day. You'd better take it for me. If it's a strike, hold your glove in place for an extra second. If it's a ball, throw it right back. And please, don't yell."

This reliance on the catcher could work because if Luciano ever suspected that he was being taken advantage of, he would have many opportunities to retaliate.

> No one I worked with ever took advantage of the situation, and no hitter ever figured out what I was doing. And only once, when Ed Herrman was calling the pitches, did a pitcher ever complain about a call. I smiled; I laughed; but I didn't say a word. I was tempted, though, I was really tempted. (Luciano and Fisher 1982, p. 166)

Ordinary business transactions are also based upon the idea that a continuing relationship allows cooperation to develop without the assistance of a central authority. Even though the courts do provide a central authority for the resolution of business disputes, this authority is usually not

invoked. A common business attitude is expressed by a purchasing agent who said that "if something comes up you get the other man on the telephone and deal with the problem. You don't read legalistic contract clauses at each other if you ever want to do business again" (Macaulay 1963, p. 61). This attitude is so well established that when a large manufacturer of packaging materials inspected its records it found that it had failed to create legally binding contracts in two-thirds of the orders from its customers (Macaulay 1963). The fairness of the transactions is guaranteed not by the threat of a legal suit, but rather by the anticipation of mutually rewarding transactions in the future.

It is precisely when this anticipation of future interaction breaks down that an external authority is invoked. According to Macaulay, perhaps the most common type of business contracts case fought all the way to the appellate courts is an action for a wrongful termination of a dealer's franchise by a parent company. This pattern of conflict makes sense because once a franchise is ended, there is no prospect for further mutually rewarding transactions between the franchiser and the parent company. Cooperation ends, and costly court battles are often the result.

In other contexts, mutually rewarding relations become so commonplace that the separate identities of the participants can become blurred. For example, Lloyd's of London began as a small group of independent insurance brokers. Since the insurance of a ship and its cargo would be a large undertaking for one dealer, several brokers frequently made trades with each other to pool their risks. The frequency of the interactions was so great that the underwriters gradually developed into a federated organization with a formal structure of its own.

The importance of future interactions can provide a guide to the design of institutions. To help promote cooperation among members of an organization, relationships should be structured so that there are frequent and durable interactions among specific individuals. Corporations and bureaucracies are often structured in just this way, as discussed in chapter 8.

Sometimes the problem is one of retarding rather than promoting cooperation. An example is the prevention of collusive business practices by avoiding the very conditions which would promote cooperation. Unfortunately, the very ease with which cooperation can evolve even among egoists suggests that the prevention of collusion is not an easy task. Cooperation certainly does not require formal agreements or even face-to-face negotiations. The fact that cooperation based upon reciprocity can emerge and prove stable suggests that antitrust activities should pay more attention to preventing the conditions that foster collusion than to searching for secret meetings among executives of competing firms.

Consider, for example, the practice of the government selecting two companies for competitive development contracts for a new military airplane. Since aerospace companies specialize to some degree in planes for either the Air Force or the Navy, there is a tendency for firms with the same specialty to face each other in the final competition (Art 1968). This frequency of interaction between two given companies makes tacit collusion relatively easy to achieve. To make tacit collusion more difficult, the government should seek methods of reducing specialization or compensating for its effects. Pairs of companies which shared a specialization would then expect to interact less often in the final competitions. This would cause later interactions between them to be worth relatively less, reducing the shadow of the future. If the next expected interaction is sufficiently far off, reciprocal cooperation in the form of tacit collusion ceases to be a stable policy.

The potential for attaining cooperation without formal agreements has its bright side in other contexts. For example, it means that cooperation on the control of the arms race does not have to be sought entirely through the formal mechanism of negotiated treaties. Arms control could also evolve tacitly. Certainly, the fact that the United States and the Soviet Union know that they will both be dealing with each other for a very long time should help establish the necessary conditions. The leaders may not like each other, but neither did the soldiers in World War I who learned to live and let live.

Occasionally a political leader gets the idea that cooperation with another major power should not be sought because a better plan would be to drive them into bankruptcy. This is an extraordinarily risky enterprise because the target need not limit its response to the withholding of normal cooperation, but would also have a strong incentive to escalate the conflict before it was irreversibly weakened. Japan's desperate gamble at Pearl Harbor, for example, was a response to powerful American economic sanctions aimed at stopping Japanese intervention in China (Ike 1967; Hosoya 1968). Rather than give up what it regarded as a vital sphere, Japan decided to attack America before becoming even further weakened. Japan understood that America was much more powerful, but decided that the cumulative effects of the sanctions made it better to attack rather than to wait for the situation to get even more desperate.

Trying to drive someone bankrupt changes the time perspective of the participants by placing the future of the interaction very much in doubt. And without the shadow of the future, cooperation becomes impossible to sustain. Thus, the role of time perspectives is critical in the maintenance of cooperation. When the interaction is likely to continue for a long time, and the players care enough about their future together, the conditions are ripe for the emergence and maintenance of cooperation.

The foundation of cooperation is not really trust, but the durability of the relationship. When the conditions are right, the players can come to cooperate with each other through trial-and-error learning about possibilities for mutual rewards, through imitation of other successful players, or even through a blind process of selection of the more successful strategies with a weeding out of the less successful ones. Whether the players trust each other or not is less important in the long run than whether the conditions are ripe for them to build a stable pattern of cooperation with each other.

Just as the future is important for the establishment of the conditions for cooperation, the past is important for the monitoring of actual behavior. It is essential that the players are able to observe and respond to each other's prior choices. Without this ability to use the past, defections could not be punished, and the incentive to cooperate would disappear.

Fortunately, the ability to monitor the prior behavior of the other player does not have to be perfect. The Computer Tournament for the Prisoner's Dilemma assumed perfect knowledge of the other player's prior choices. In many settings, however, a player may occasionally misperceive the choice made by the other. A defection may go undetected, or a cooperation may be misinterpreted as a defection. To explore the implications of misperception, the

first round of the tournament was run again with the modification that every choice had a 1 percent chance of being misperceived by the other player. As expected, these misunderstandings resulted in a good deal more defection between the players. A surprise was that TIT FOR TAT was still the best decision rule. Although it got into a lot of trouble when a single misunderstanding led to a long echo of alternating retaliations, it could often end the echo with another misperception. Many other rules were less forgiving, so that once they got into trouble, they less often got out of it. TIT FOR TAT did well in the face of misperception of the past because it could readily forgive and thereby have a chance to reestablish mutual cooperation.

The role of time perspective has important implications for the design of institutions. In large organizations, such as business corporations and governmental bureaucracies, executives are often transferred from one position to another approximately every two years.[2] This gives executives a strong incentive to do well in the short run, regardless of the consequences for the organization in the long run. They know that soon they will be in some other position, and the consequences of their choices in their previous post are not likely to be attributed to them after they have left their position. This gives two executives a mutual incentive to defect when either of their terms is drawing to an end. The result of rapid turnover could therefore be a lessening of cooperation within the organization.

As pointed out in chapter 3, a similar problem arises when a political leader appears to have little chance of reelection. The problem becomes even more acute with a lame duck. From the point of view of the public, a politician facing an end of career can be dangerous because of the increased temptation to seek private goals rather than maintain a pattern of cooperation with the electorate for the attainment of mutually rewarding goals.

Since the turnover of political leaders is a necessary part of democratic control, the problem must be solved another way. Here, political parties are useful because they can be held accountable by the public for the acts of their elected members. The voters and the parties are in a long-term relationship, and this gives the parties an incentive to select candidates who will not abuse their responsibilities. And if a leader is discovered giving in to temptation, the voters can take this into account in evaluating the other candidates of the same party in the next election. The punishment of the Republican party by the electorate after Watergate shows that parties are indeed held responsible for the defections of their leaders.

In general, the institutional solutions to turnover need to involve accountability beyond the individual's term in a particular position. In an organizational or business setting, the best way to secure this accountability would be to keep track not only of the person's success in that position, but also the state in which the position was left to the next occupant. For example, if an executive sought a quick gain by double-crossing a colleague just before transferring to a new plant, this fact should be taken into account in evaluating that executive's performance.

Cooperation Theory has implications for individual choice as well as for the design of institutions. Speaking personally, one of my biggest surprises in working on this project has been the value of provocability. I came to this project believing one should be slow to anger. The results of the Computer Tournament for the Prisoner's Dilemma demonstrate that it is actually better to respond quickly to a provocation. It turns out that if one waits to respond to

uncalled for defections, there is a risk of sending the wrong signal. The longer defections are allowed to go unchallenged, the more likely it is that the other player will draw the conclusion that defection can pay. And the more strongly this pattern is established, the harder it will be to break it. The implication is that it is better to be provocable sooner, rather than later. The success of TIT FOR TAT certainly illustrates this point. By responding right away, it gives the quickest possible feedback that a defection will not pay.

The response to potential violations of arms control agreements illustrates this point. The Soviet Union has occasionally taken steps which appear to be designed to probe the limits of its agreements with the United States. The sooner the United States detects and responds to these Soviet probes, the better. Waiting for them to accumulate only risks the need for a response so large as to evoke yet more trouble.

The speed of response depends upon the time required to detect a given choice by the other player. The shorter this time is, the more stable cooperation can be. A rapid detection means that the next move in the interaction comes quickly, thereby increasing the shadow of the future as represented by the parameter w. For this reason the only arms control agreements which can be stable are those whose violations can be detected soon enough. The critical requirement is that violations can be detected before they can accumulate to such an extent that the victim's provocability is no longer enough to prevent the challenger from having an incentive to defect.

The tournament results concerning the value of provocability are complemented by the theoretical analysis of what it takes for a nice rule to be collectively stable. In order for a nice rule to be able to resist invasion, the rule must be provocable by the very first defection of the other player (proposition 4 in chapter 3). Theoretically, the response need not come immediately, and it need not occur with certainty, but it must have a real probability of coming eventually. The important thing is that the other player does not wind up having an incentive to defect.

Of course, provocability has a danger. The danger is that if the other player does try a defection, retaliation will lead to further retaliation, and the conflict will degenerate into an unending string of mutual defections. This can certainly be a serious problem. For example, in many cultures blood feuds between clans can continue undiminished for years and even generations (Black-Michaud 1975).

This continuation of the conflict is due to the echo effect: each side responds to the other's last defection with a new defection of its own. One solution is to find a central authority to police both sides, imposing a rule of law. Unfortunately this solution is often not available. And even when there is a rule of law, the costs of using the courts for routine affairs such as enforcement of business contracts can be prohibitive. When the use of a central authority is impossible or too expensive, the best method is to rely on a strategy which will be self-policing.

Such a self-policing strategy must be provocable, but the response must not be too great lest it lead to an unending echo of defections. For example, suppose that the Soviet Union in conjunction with the other Warsaw Pact countries undertakes a partial mobilization of its armed forces throughout Eastern Europe. This mobilization would give the Soviets an added advantage if conventional war were to break out. A useful response from NATO would be to increase its own state of alert. If additional troops moved

from the Soviet Union to Eastern Europe, NATO should respond with additional troops moved from the United States. Betts (1982, pp. 293–94) recommends that this type of response be automatic so that it can be made clear to the Soviets that such increases in NATO readiness are standard procedure and take place only after Soviet mobilization. He also recommends that the response be limited, say one American division moved for every three Soviet divisions mobilized. In effect, this would help limit the echo effects.

Limited provocability is a useful feature of a strategy designed to achieve stable cooperation. While TIT FOR TAT responds with an amount of defection exactly equal to the other's defection, in many circumstances the stability of cooperation would be enhanced if the response were slightly less than the provocation. Otherwise, it would be all too easy to get into a rut of unending responses to each other's last defection. There are several ways for an echo effect to be controlled. One way is for the player who first defected to realize that the other's response need not call for yet another defection. For example, the Soviets might realize that NATO's mobilization was merely a response to their own, and hence need not be regarded as threatening. Of course the Soviets might not see it that way, even if the NATO response was automatic and predictable. Therefore, it is also useful if the NATO response is somewhat less than proportional to the Soviet mobilization. Then if the Soviet response is also somewhat less than the NATO mobilization, the escalation of preparations can become stabilized, and then possibly reversed for a return to normal.

Fortunately, friendship is not necessary for cooperation to evolve. As the trench warfare example demonstrates, even antagonists can learn to develop cooperation based upon reciprocity. The requirement for the relationship is not friendship, but durability. The good thing about international relations is that the major powers can be quite certain they will be interacting with each other year after year. Their relationship may not always be mutually rewarding, but it *is* durable. Therefore, next year's interactions should cast a large shadow on this year's choices, and cooperation has a good chance to evolve eventually.

Foresight is not necessary either, as the biological examples demonstrate. But without foresight, the evolutionary process can take a very long time. Fortunately, humans do have foresight and use it to speed up what would otherwise be a blind process of evolution. The most striking example of this was the difference between the first and second rounds of the Computer Prisoner's Dilemma Tournament. In the first round the contestants were professional game theorists who represented the state of the art in the understanding of how to do well in the iterated Prisoner's Dilemma. When their rules were paired with each other, the result was an average score per move of 2.10 which is only slightly better than halfway from $P=1$ (the punishment for mutual defection) to $R=3$ (the reward for mutual cooperation). The players in the second round did much better, scoring 2.60, which is a little better than three-quarters of the way from the mutual punishment to the mutual reward.[3] Thus, the players were able to use the results of the first round, including the value of reciprocity, to anticipate what would work well in the second round. On the whole, their foresight paid off with substantially higher scores.

The result was that the second round was more sophisticated than the first. Cooperation based upon reciprocity was firmly established. The various attempts at exploitation of the unsophisticated entries of the first round all failed in

the environment of the second round, demonstrating that the reciprocity of strategies like TIT FOR TAT is extraordinarily robust. Perhaps it is not too much to hope that people can use the surrogate experience of the Computer Tournament to learn the value of reciprocity for their own Prisoner's Dilemma interactions.

Once the word gets out that reciprocity works, it becomes the thing to do. If you expect others to reciprocate your defections as well as your cooperations, you will be wise to avoid starting any trouble. Moreover, you will be wise to defect after someone else defects, showing that you will not be exploited. Thus you too will be wise to use a strategy based upon reciprocity. So will everyone else. In this manner the appreciation of the value of reciprocity becomes self-reinforcing. Once it gets going, it gets stronger and stronger.

This is the essence of the ratchet effect which was established in chapter 3: once cooperation based upon reciprocity gets established in a population, it cannot be overcome even by a cluster of individuals who try to exploit the others. The establishment of stable cooperation can take a long time if it is based upon blind forces of evolution, or it can happen rather quickly if its operation can be appreciated by intelligent players. The empirical and theoretical results of this book might help people see more clearly the opportunities for reciprocity latent in their world. Knowing the concepts that accounted for the results of the two rounds of the Computer Prisoner's Dilemma Tournament, and knowing the reasons and conditions for the success of reciprocity, might provide some additional foresight.

We might come to see more clearly that there is a lesson in the fact that TIT FOR TAT succeeds without doing better than anyone with whom it interacts. It succeeds by eliciting cooperation from others, not by defeating them. We are used to thinking about competitions in which there is only one winner, competitions such as football or chess. But the world is rarely like that. In a vast range of situations mutual cooperation can be better for *both* sides than mutual defection. The key to doing well lies not in overcoming others, but in eliciting their cooperation.

Today, the most important problems facing humanity are in the arena of international relations, where independent, egoistic nations face each other in a state of near anarchy. Many of these problems take the form of an iterated Prisoner's Dilemma. Examples can include arms races, nuclear proliferation, crisis bargaining, and military escalation. Of course, a realistic understanding of these problems would have to take into account many factors not incorporated into the simple Prisoner's Dilemma formulation, such as ideology, bureaucratic politics, commitments, coalitions, mediation, and leadership. Nevertheless, we can use all the insights we can get.

Robert Gilpin (1981, p. 205) points out that from the ancient Greeks to contemporary scholarship all political theory addresses one fundamental question: "How can the human race, whether for selfish or more cosmopolitan ends, understand and control the seemingly blind forces of history?" In the contemporary world this question has become especially acute because of the development of nuclear weapons.

The advice in chapter 6 to players of the Prisoner's Dilemma might serve as good advice to national leaders as well: don't be envious, don't be the first to defect, reciprocate both cooperation and defection, and don't be too clever. Likewise, the techniques discussed in chapter 7 for promoting cooperation in the Prisoner's Dilemma might also be useful in promoting cooperation in international politics.

The core of the problem of how to achieve rewards from cooperation is that trial and error in learning is slow and painful. The conditions may all be favorable for long-run developments, but we may not have the time to wait for blind processes to move us slowly toward mutually rewarding strategies based upon reciprocity. Perhaps if we understand the process better, we can use our foresight to speed up the evolution of cooperation.