

On Six Advances in Cooperation Theory

Robert Axelrod

School of Public Policy

University of Michigan

Ann Arbor, MI 48109, USA

Axe@umich.edu

January 2000

Prepared for a Special Issue of *Analyse & Kritik* on

The Evolution of Cooperation

Abstract: The symposium included in this issue of *Analyse & Kritik* extends the basis of Cooperation Theory as set forth in Axelrod's *Evolution of Cooperation* (1984). This essay begins with an overview of Cooperation Theory in terms of the questions it asks, its relationship to game theory and rationality, and the principal methodologies used, namely deduction and simulation. This essay then addresses the issues raised in the symposium, including the consequences of extending the original paradigm of the two person iterated Prisoner's Dilemma to take into account such factors as nonsimultaneous play, the ability to offer hostages for performance, social networks of interaction, information sharing that can support reputations, learning behavior, envy, misunderstanding, and an option to exit. The essay places the contributions of this symposium in the context of previous research on these and related issues.

Introduction

I am most grateful to Professor Baurmann for organizing this symposium and for offering me the opportunity to respond to the very interesting papers included here.¹ There is nothing more gratifying for a scholar than to see ones work used by others as a foundation for creative and productive advances. For this reason, I am also grateful to the authors of these papers, from whom I have learned a great deal.

The range of extensions and issues is truly impressive, including information sharing that can support reputations, nonsimultaneous play, the ability to offer hostages for performance, social networks of interaction, learning behavior, envy, misunderstanding, and an option to exit. The six papers in this symposium present advances on all of these fronts. The principal role of this essay is to place this work in the context of previous research on these and related issues. The last comprehensive review was a dozen years ago (Axelrod and Dion, 1988). Citations to *The Evolution of Cooperation* (Axelrod 1984) are now growing at the rate of over 300 per year. The literature on Cooperation Theory is now so large that the authors can be forgiven for not being fully cognizant of all work related to their own research topics. I hope that placing the contributions of this symposium in the context of recent work in the field will accomplish two things. First, it can help lead to a deeper appreciation of what has been established so far. Second, placing the present work in context of related work will help identify some promising opportunities for further advances.

This essay begins with an overview of Cooperation Theory in terms of the questions it asks, its relationship to game theory and rationality, and the principal methodologies used, namely deduction and simulation.

The basic problem that Cooperation Theory addresses is the common tension between what is good for the individual actor in the short run, and what is good for the group in the long run. The Prisoner's Dilemma embodies this tension in a particularly simple and compelling manner. For that reason, the Prisoner's Dilemma has become the foundation for most work in Cooperation Theory, across a wide range of disciplines. But

¹ For financial support in preparing this response, I thank the Intel Corporation and the

as we shall see, there are other games that are useful for studying aspects of the fundamental problem of cooperation that are not captured by the standard Prisoner's Dilemma.

Regardless of the theoretical details, however, virtually all of Cooperation Theory employs game theory as the basis for analysis. Game theory begins with a set of actors, each of whom has a set of choices. When the players each make their choice, there is an outcome that is jointly determined by the choices of the players. The outcome determines the payoffs to the players. Consider the one-move two-person Prisoner's Dilemma, as an example of a game. The choices are cooperate or defect, resulting in four possible outcomes. The possible payoffs are the reward for mutual cooperation, R , which is greater than the punishment for mutual defection, P . The dilemma is caused by the fact that the temptation payoff for unilateral defection, T , is greater than the sucker's payoff for unilateral cooperation, S . The Prisoner's Dilemma is defined by $T > R > P > S$. A second condition is usually added so that mutual cooperation is better than coordinated alternation of cooperation: $R > (S+T)/2$.

In an iterated game, a player can use a strategy that relies on the information available so far to decide at each move which choice to make. Since the players do not know when the game will end, they both have an incentive and an opportunity to develop cooperation based upon reciprocity. The shadow of the future provides the basis for cooperation, even among egoists. An example of a reciprocating strategy for the iterated Prisoner's Dilemma is Tit for Tat which cooperates on the first move, and then does whatever the other player did on the previous move.

Game theory allows a very rich way of analyzing what will happen in a specific strategic context. To specify a game, one needs to specify the players, the choices, the outcomes as determined jointly by the choices, and the payoffs to the players associated with the outcomes. One more thing is needed. What is needed is a way of determining how the players will make their choices, or in the case of an iterated game how they will select their strategies. Traditionally, game theory has calculated what players will do by assuming the players are rational, that they know the other players are rational, and that everyone has the ability to do unlimited calculation. Clearly, the assumption of rationality is very strong.

The rationality assumption of traditional game theory has been widely challenged. Among the leaders of the challenge is Herbert Simon (1982), who has emphasized that people have limited knowledge of their situations, limited ability to process information, and limited time to make choices. People are therefore likely to use rules of thumb rather than detailed calculation, more likely to experiment than try to determine an optimal response, and more likely to imitate someone who seems to be doing well rather than rely completely on their own experience (March 1978). Cooperation Theory has taken these observations seriously, and is as likely to study adaptive actors as it is to study fully rational actors. It should be noted that in recent years, game theory as a whole has begun to relax the assumption of rational actors, and studied various forms of adaptive behavior (Samuelson 1997, Hofbauer and Sigmund 1998, Fudenberg and Levine 1998, Young 1998). The emphasis on adaptive actors and evolutionary processes that has characterized Cooperation Theory from the beginning is now becoming quite widespread throughout game theory.

Cooperation Theory has three central theoretical questions.

1. Under what conditions can cooperation emerge and be sustained among actors who are egoists?
2. What advice can be offered to a player in a given setting about the best strategy to use?
3. What advice can be offered to reformers who want to alter the very terms of the interaction so as to promote the emergence of cooperation?

The papers in this symposium address all three of these theoretical issues.² Two papers in this symposium study issues that arise in the original strategic setting of the iterated Prisoner's Dilemma, while four papers analyze the consequences of making certain modifications in that setting. All the papers undertake their strategic analysis within the general framework of game theory in general, and Cooperation Theory in particular.

The extreme simplicity of the Prisoner's Dilemma paradigm proved to have several important benefits over the years. First, it allowed a set of theorems to be proved about the conditions under which cooperation can get started and be sustained (e.g., Axelrod and Dion 1988, Bendor and Swistak 1997). Second, it allowed both professional game theorists and amateur computer hobbyists to devise an impressive range of more or less sophisticated strategies with which to play the game. These strategies provided the basis for two computer tournaments, which in turn provided powerful evidence about the performance and robust success of the strategy of Tit for Tat (Axelrod 1984). Third,

² Cooperation Theory also addresses empirical questions about the accuracy of the predictions derived from the theory, and about the extent to which the dynamics of

these results have inspired a good deal of empirical work demonstrating that cooperation based upon reciprocity does indeed exist between individuals, between nations, and even among animals.³

The extreme simplicity of the Prisoner's Dilemma paradigm has allowed the authors of this symposium to extend the basic framework without getting too complicated. The authors adhere to the KISS principle of the old army slogan, "Keep it simple, stupid"(Axelrod 1997a, p. 5). The KISS principle is vital because of the character of the research community. When surprising results are discovered – as they often are in this symposium – it is very helpful to be confident that we can understand everything that went into the model that produced the surprises.

Before turning to the specifics of the individual papers, there are two questions that are relevant to several of them that can best be addressed at the start. These questions are:

1. What are the relative advantages and disadvantages of studying social processes with computer simulation compared to the more established method of deductive reasoning?

2. How should we regard the strategy of permanent retaliation (the so-called "Grim Trigger") which is used for analytic purposes in two of the papers?

historical cases are illuminated by the theory (e.g., Axelrod 1984, Axelrod and Dion 1988, Axelrod 1997a).

³ For examples from fish to nations, see the citations in Axelrod and Dion (1988). Recent evidence suggests the Prisoner's Dilemma exists even for a virus (Nowak and Sigmund 1999, Turner and Chao 1999). In addition to reciprocity based on the shadow of the future, other factors that tend to support cooperation are relatedness of the players (Hamilton 1964, Dawkins 1989) and internalization of social norms (Simon 1990).

The papers in this symposium use two basic techniques to generate results from models: deduction and simulation. These two techniques have complementary advantages and disadvantages. Deduction involves specifying a set of axioms, and proving theorems based on them. Simulation also involves specifying a set of assumptions, but instead of proving theorems, it works by generating “histories” and then analyzing patterns in those histories. Deduction has several advantages over simulation. First, any theorem that can be proved is definitely true. One’s confidence in a theorem is complete. In contrast, the detection of a pattern in simulated data is typically characterized by some degree of confidence. A typical statement about a statistical pattern is that there is less than 5% chance that it would have been caused by a mechanism that generated data at random. Clearly, certainty is better than likelihood. The other advantage of deduction is that a theorem typically reveals the role of parameters, whereas simulation has to rely on trying out specific values of the parameters. For example, it is a theorem that if the other player is using Tit for Tat in an iterated Prisoner’s Dilemma, a player can do no better than using Tit for Tat when $w > \max((T-R)/(T-P), (T-R)/(R-S))$ where w is the discount rate per move (Axelrod 1981). Once this theorem is established, the implications for any combination of parameters for the payoffs and the discount rate can be immediately established. In a simulation, on the other hand, the analysis would have to be repeated for many combinations of the parameters to see their combined effects. And even after doing many simulation runs, one might not be sure that there would be some unexplored combination of the parameters that might lead to a different result. So to the extent that the desired results can be attained by deduction, simulation is a second-best technique.

What gives simulation its power is that it can often be used when deduction is not possible. Even simple models often involve effects that are difficult or even impossible to pin down by deduction. This is especially likely to be the case when there are many elements in the system which interact in non-linear ways. In fact, this is exactly the case in many problems that Cooperation Theory is meant to address. There is often a whole population of agents, and they each interact with many others. The results might well depend on the emerging pattern of interaction, as well as what the agents private experience as they go. Existing mathematics may simply be inadequate to predict or account for the resulting histories. The power of simulation is that histories can be generated once one specifies the assumptions underlying the dynamics of the model. For example, if the strategies and interaction rules are specified, then a simulation can generate histories that follow those rules. Patterns can then be discovered by examining populations of histories, each of which consists of a population of agents. But as Buskens and Weesie (2000) point out, it pays to be cautious about generalizing from simulation results since until a firm analytic understanding is achieved, one can not be completely confident how well the results of particular simulation runs will generalize to other conditions.

Simulation has proven especially useful in the study of adaptive agents. This is because adaptive agents typically update their strategies based on experience in ways that might be easy to specify, but hard to analyze mathematically. This is the reason that Hegselmann and Flache (2000) use simulation in the comparative analysis of rational and adaptive agents. Similarly the Edk-Group (2000) uses simulation to study agents with limited memory. In this study, there is also evolutionary turnover of agents in the

population, with the less successful agents being replaced by agents using more successful strategies. While it is possible to get analytic results with some evolutionary models, these models tend to become intractable fairly quickly when mutation is allowed. Therefore, simulation has been the preferred method for treating evolutionary models with mutation.

Simulation is a way of doing thought experiments (Axelrod 1997b). While the assumptions may be simple, the consequences may not be obvious at all. The large-scale effects of locally interacting agents often yield what are known as “emergent properties” of the system. Emergent properties are often surprising because it can be hard to predict the full consequences of even simple forms of interaction. A good example is Schelling's (1978) model of residential tipping. In this model a family moves only if more than one third of its immediate neighbors are of a different type. The result is that very segregated neighborhoods form even though everyone is initially placed at random, and everyone is somewhat tolerant. It would be difficult to establish this result by deduction. But simulation demonstrates the result clearly and compellingly. Put another way, simulation provides an existence proof that certain results are possible from a given set of assumptions. In this symposium, as in most game theory, both deduction and simulation aim more for the illumination of basic principles than for accurate representation of any particular realistic application. The goal is to enrich our understanding of fundamental processes that may appear in a variety of applications.⁴

A good example of the difference between deductive and simulation approaches is provided by the analysis of the merits of a particular strategy for the iterated Prisoner's

Dilemma used in some variation in two of the papers. This is the strategy known as Grim Trigger or Permanent Retaliation. It starts by cooperating, and continues to cooperate until the other player's first defection; then it never cooperates again. The Grim Trigger strategy imposes the most severe punishment available for the smallest departure from cooperation, namely a response of eternal detection (Friedman 1971). As Hegselmann and Flache (2000) point out, it can be proven that the conditions to sustain cooperation with Grim Trigger are necessary conditions for the possibility of any form of conditional cooperation. Put another way, Grim Trigger can sustain cooperation in the iterated Prisoner's Dilemma under the least favorable circumstances of any strategy that can sustain cooperation. In a variant of the Prisoner's Dilemma designed to study trust, Buskens and Weesie (2000) used the analogy of Grim Trigger as the strategy that starts out trusting other players, but *never* again trusts a player when there is information that that player abused anyone's trust. As Buskens and Weesie (2000) point out, threatening "eternal" punishment is the most effective way to sustain trust because the other player's loss is maximized after trust is abused even once. By assuming that the basic strategy of the trust game is Grim Trigger, Buskens and Weesie (2000) are able to prove a whole series of quite general theorems about when trust can be sustained.

While Grim Trigger allows the deduction about the minimal conditions which are needed to sustain cooperation (or trust), simulation helps to show that Grim Trigger is actually a very dangerous strategy for the user, as well as for the other player. Consider the experience of the two rounds of the Prisoner's Dilemma computer tournament. In both rounds, Professor Friedman submitted Grim Trigger as his entry. In the first round,

⁴ Modeling can be used for other purposes as well. These include prediction, performance

it scored 7th out of 14 submitted entries (Axelrod 1984, p. 193). In the second round, it scored 52nd out of 62 submitted entries (Axelrod 1984, p. 195). Clearly, it was not a very successful strategy. What success it did have was due to the fact that it was never the first to defect. Being a nice strategy in this sense meant that it did as well as possible with the other nice strategies in the tournaments. In fact, being nice was the single best predictor of how well a strategy did in the tournaments. But other than being nice, Grim Trigger did not have much going for it. In fact, of the 39 nice strategies in the second round, Grim Trigger did worse of all.

The problem of course is that if the other player ever defected, Grim Trigger never cooperated again. Unending defection is a good way to play with completely uncooperative strategies, but it is not a good way to play with responsive strategies that might be trying an occasional defection to see what they can get away with. Typically, the unending string of defections from Grim Trigger led the exploratory player to sooner or later simply give up and defect almost all of the rest of the game. This resulted in low scores for both Grim Trigger and the exploratory player. Note that in the tournaments, Grim Trigger was not able to communicate its threat of massive retaliation in advance. Once the other player provoked Grim Trigger, it was too late.

Another problem with Grim Trigger is that it is highly susceptible to noise. If it mistakenly believes that the other player defected, it will never cooperate again. Just as a little exploratory behavior by the other player can set off Grim Trigger, so can a little noise. And once set off, Grim Trigger not only punishes the other player but also itself

of tasks, training, entertainment, and education (Axelrod, 1997b).

suffers from the other player's retaliation. And the other player's retaliation typically becomes almost as consistent as Grim Trigger's behavior. Thus both players suffer.

A comparison with Tit for Tat shows that both are nice strategies, being never the first to defect. And both are provokable by the first defection of the other. But the difference is that Tit for Tat is completely forgiving after one punishment for one defection, while Grim Trigger is completely unforgiving after one defection and provides maximal punishment for even a single defection.

In sum, Grim Trigger seems like a good idea, but isn't. It does offer the maximal incentive for the other player to completely avoid defection.⁵ But if the other player doesn't know that it is facing Grim Trigger, it can't adjust its behavior until it is too late. Any experimentation (or noise) will end in trouble for both sides. Thus in a world of more or less sophisticated players where you can observe the other's behavior but can not know its strategy in advance, Grim Trigger is likely to be a poor performer.

This discussion illustrates two principles. First, what makes good advice depends not only on the deduced properties of the strategy in question, but also on the exact conditions under which the strategy will be used. In a world of adaptive agents, even a fully rational player needs to take into account that the other players are likely to be experimenting rather than optimizing. Second, simulations offer a rich possibility for checking the effectiveness of strategic ideas in environments that are highly diverse.

⁵ Another potential advantage of Grim Trigger is that it can exploit strategies which never give up trying to cooperate, even after being repeatedly punished. Linster (1992) shows how this can happen. Linster's simulation uses an environment composed of two-state Moore machines. This environment provides just the kind of strategies that Grim Trigger is good at exploiting.

Having considered the questions related to the symposium as a whole, we are now ready to turn to the individual papers.

Timing of Choices

Abell and Reyniers (2000) extend the basic paradigm of the Prisoner's Dilemma by considering what happens when the players do not necessarily make their choices at the same time. They then study the process of generalized reciprocity that can arise in this setting. They consider three-player games, as well as two player games in order to capture the idea that actor 1 may help actor 2 now in the expectation that actor 2 or someone else (actor 3) will reciprocate later when actor 1 needs help. This is an important extension of the original paradigm because it allows the analysis of certain settings that are not well represented in the original paradigm. The paper provides an interesting and useful set of deductive results focusing on the conditions that are required for cooperation to be sustained in such a setting.

Let me take this opportunity to place this work in a broader context. We now have three ways to model the sequencing of moves between two players:

a. The (standard) iterated Prisoner's Dilemma in which the two players move at the same time, and then make their next move after learning what the other player did on the previous move.

b. The alternating Prisoner's Dilemma in which the players take turns. The leader moves first, and the follower moves next, then the leader moves again, and so on (Nowak and Sigmund 1994).

c. The bilateral Prisoner's Dilemma in which either, neither or both have an opportunity to help the other in each round. Player 1 has a certain probability of moving in each round, and player 2 has an independent and possibly different probability of moving in each round (Abell and Reyniers 2000).

We can ask, "For a given application, which of the settings is the best model?" Here is my answer.

a. The standard game with its simultaneous moves corresponds to situations in which each player gets to move at every opportunity. The length of time between moves might be due exogenous circumstances such as when the two players happen to meet each other. Or the players can be in continual contact, and the length of time between moves can correspond to the time it takes either of them to learn what the other did and implement a new choice. For example, if two nations are in an arms race or two companies are in a price competition, then the time between moves corresponds to the time it takes a player to observe a change in the other's behavior, and implement a new choice in response. This might be an annual arms budget cycle, or a weekly price setting cycle.

b. The alternating game corresponds to situations in which the players take turns because they both can not receive help at the same time. A good biological example is young male baboons who alternate the role of distracting the attention of the dominant male while the other has the opportunity to mate with an estrous female (Parker 1977; Trivers 1971; cited in Hauert and Schuster 1998). In human situations, one person might receive help one day while asking to receive help the next day. In order for the game to be strictly alternating, the opportunity to give and receive help must switch each time.

While the order of moves might be controlled by some outside circumstance or authority, a common reason players alternate is that they keep track of whose turn it is receive help.

c. The bilateral game means that either or both players might have an opportunity to help the other. This situation would arise when the opportunities themselves are beyond the control of the players, and the opportunity for one player to help is independent of the opportunity for the other player to help. Abell and Reyniers (2000) do not give any specific examples. One could imagine, however, two students studying for an exam who might have some things they need help with. The first student might understand some things the second student doesn't. The second might understand some things the first student doesn't. Or both. Or neither. Thus the bilateral game studied by Abell and Reyniers corresponds to a situation in which at any point in time the players might need help and be able to offer help, and that this occurs in a strictly uncorrelated manner.

d. The previous two settings suggest a fourth possibility, that I would call the single resource game. This is the case where only one player at a time can get help, and the need is determined exogenously. For example, if you need a loan, I might lend you money in the expectation that someday I might need a loan. Unlike the alternating case, we don't necessarily take turns since we can't control when we might need help. Unlike the bilateral case, we can't both help each other at the same time. The single resource could be food or money, or anything else that has uncertain availability and diminishing marginal returns. The diminishing marginal returns assumption guarantees that person would be happy to offer some of the resource in times of plenty provided there was

sufficient chance of getting enough back in hard times to make the interaction worthwhile.

In sum, which game is most appropriate depends on the relationship between the players. The standard Prisoner's Dilemma is the appropriate model when players can always help each other. The alternating game is appropriate when the players can or must take turns. The bilateral game is appropriate when opportunities for help are exogenous and independent. The single resource game is appropriate when there is a single resource that one player may be able to offer the other, but the opportunities do not necessarily alternate.

The standard Prisoner's Dilemma has a huge literature. The alternating game has developed a substantial literature over just the last few years.⁶ The bilateral game (and its three-person version for generalized reciprocity) is just beginning with Abell and Reyniers (2000). To my knowledge, no one has systematically analyzed the single resource game in these terms (with exogenous but nonsimultaneous needs).

What difference does the setting make? William Hamilton and I made the claim (Axelrod and Hamilton 1981) noted by Abell and Reyniers that it would make little difference if the moves were sequential rather than simultaneous. We didn't specify what we meant by "sequential" or what we meant by "make little difference." Now that three different ways that the moves can be sequential have been identified, one could begin to sort out the answer. A complete assessment is beyond the scope of this essay.

⁶ For theoretical treatments of the alternating Prisoner's Dilemma see Nowak and Sigmund (1994), Frean (1994); Wedekind and Milinski (1996); Boerlijst, Nowak and Sigmund (1997); Leimar (1997); and Hauert and Schuster (1998). Some of these papers find merit not only in reciprocating strategies, but also in Pavlovian strategies. Wedekind

Nevertheless, based on the literature so far, it still seems reasonable to suppose that the main conclusion of the basic paradigm still holds: cooperation based on reciprocity can be sustained if and only if the payoff parameters and the shadow of future are favorable enough. In all cases, a key role is played by the shadow of the future that is interpreted as the probability the game will end at a given time, or the discount rate between moves (Axelrod 1984). In all cases, the best strategy to use depends in part on the strategy the other player may be using. If the other player is likely to be sufficiently responsive, and the payoffs and shadow of the future are sufficiently favorable, recommending a reciprocal strategy still seems like robust advice.

Hostages

Raub and Weesie (2000) consider a different way to promote cooperation. Instead of iterating the game, they analyze the possibility that a player (called the trustee) can voluntarily provide a hostage, such as a bond. The hostage is intended to convince the other player (called the trustor) that the trustee will in fact cooperate. They show how this can work to promote cooperation even in one-sided Prisoner's Dilemma where the trustee moves just once, and then the trustor responds just once. The paper demonstrates how hostages help promote trust in three different ways: reducing the incentive of the trustee to abuse the trust, reducing the cost to the trustor if the trust is abused, and serving

and Milinski (1996) even compare how biology students play the standard and alternating game.

as a useful signal about the characteristics of the trustee. Elucidating the role of hostage posting as a useful signal is a particularly valuable contribution.⁷

Historically, hostages have often been used to guarantee performance. The typical case was for an imperial authority or conquering power to take hostages from a village to guarantee the payment of taxes in the form of money or labor services. The Chinese used this technique as early as the fourth century BC (Dewey 1988). The Romans, the Mongols, and almost everyone else it seems also used hostage taking. Involuntary hostage taking offends our deepest sense of justice not only because it serves the interests of the conquerors, but also because it involves punishing the innocent. Indeed, a Geneva Convention has now outlawed the practice.

A historically important variant of hostage taking is the use of the entire population of a village to guarantee the performance of each of its members. The typical method was to impose taxes on a village, rather than on a household. Then if someone runs away, the rest of the village has to make up their share of the tax. This forces the village to organize itself to prevent runaways. The result is that the entire village is held hostage for the performance of each of its members. In Russia this system was introduced by the Mongols, but flourished under the Czars long after their departure (Dewey 1988). One may plausibly speculate that the long experience of coercive village responsibility may have helped shape Russian popular attitudes against individualism.

Raub and Weesie (2000) quite rightly trace the game theoretic treatment of hostages back to Schelling (1960). For Schelling, the existence of hostages was not a

⁷ An exemplary feature of this paper is the way it uses a single example (of a lawyer and a law firm) to explain and motivate a series of ever more elaborate models of the trust process.

matter of choice. Instead, the lack of defense against nuclear weapons meant that the populations of entire countries were hostages. During the Cold War, this may have led to some degree of trust that the weapons would not be used. But clearly, the degree of trust was not very great and there was always some reciprocal fear of surprise attack (Schelling 1966). During the Cuban Missile Crisis, for example, we came perilously close to major war despite the existence of hostages (Allison and Zelikow 1999). Since Schelling's time, game theory tied to empirical analysis has come a long way toward understanding strategic issues of using threats based upon hostages (e.g., Powell 1999).

Fortunately, the taking of human hostages for tax collection has become rare in modern societies. Equally fortunate, the end of the Cold War has reduced our reliance on the vulnerability of hostage populations as a means of deterring war. Raub and Weesie (2000) show how the voluntary posting of hostages in the form performance bonds can actually promote trust and cooperation. What we need now is a better understanding of the subtle relationship between voluntary posting of bonds and coerced posting. For example, if a law firm places trust in a newly hired lawyer by providing extensive training, the firm may want some guarantee that the lawyer will stay with the firm. We want to be sure that the law firm is not allowed to use coercive ways of making the lawyer post a hostage. The analysis by Raub and Weesie (2000) of how and when voluntary hostage taking works can in the future serve the additional function of helping to identify the incentives and dangers of coercive hostage posting.

Social Networks

Buskens and Weesie (2000) consider how reputation effects can promote cooperation. Building on Raub and Weesie (2000), they use a trust game, which is related to a one-sided Prisoner's Dilemma. Instead of using hostages to provide the basis of trust, this paper shows that information sharing can do the job. The specific form of the information sharing is an opportunity for communication between one trustor and another. If a trustor informs the next trustor, she communicates not only her own experiences, but also all the information she has obtained from previous trustors. This information transfer allows reputations to be established, providing incentives to cooperate even if a player may never play again with the same partner.

The paper's greatest strength is in its analysis of the role of social structure in supporting cooperation based upon reputation. While some useful analytic results are derived concerning specific types of social networks, the authors conclude that computer simulation will be needed to go further. The results obtained show that for a given social structure, the payoff parameters work as expected. In addition the social structure itself has a large effect on how favorable these parameters have to be to support cooperation based on reputation. The reason is that the spread of reputation depends heavily on who informs whom of what. When the social structure is favorable, cooperation based on reputation can be sustained even when two players may never meet again.

This work combines a concern with reputation with a concern with social structure. Previous work has mainly focused on one or the other of these factors.

Nowak and Sigmund (1998a and 1998b) study a closely related model for the spread of reputation. They also used a game related to the one-move, one-sided Prisoner's Dilemma. In their model, information about the number of times that the other

player cooperated was public knowledge. Instead of relying on a trigger strategy, as in Buskens and Weesie (2000), players had different thresholds of tolerance for the other player's past behavior. Like Buskens and Weesie (2000), Nowak and Sigmund found that cooperation based on reputation can be sustained under certain conditions, even though two players may never meet more than once. Again, the key was the spread of information that allowed reputations to be formed. It would be interesting to compare the Nowak and Sigmund model using the networks analyzed by Buskens and Weesie (2000) in order to see how robust the results about social structure are with respect to what information is shared and how it is used.

While Buskens and Weesie (2000) study the role of social structure in supporting cooperation via reputation, there is also an extensive literature on how social structure can support cooperation even without information sharing between players. The most common way this is demonstrated with the social structure of a two dimensional lattice, in which players interact only with their four immediate neighbors. These studies then assume that players update their strategy by adapting the strategy of a neighbor who did better than they did (Axelrod 1984, 158-68; Pollock 1989; Nowak and May 1992; Lindgren and Nordahl 1994; Nowak et al. 1994; Grim 1997; Nakamura, Matusda and Iwasa 1997). These studies show how a highly structured social interaction pattern can sustain cooperation in circumstances that would not have sustained it if the players mixed freely. In fact, various forms of social structure can sustain cooperation even without information sharing between players. Random networks can do the job, as long as the relationships are fixed (Cohen et al. 1999). In fact, cooperation can even be sustained

when the basis of the social structure is merely a tendency to interact with others who are similar on a completely arbitrary property (Riolo 1997, Cohen et al. 1999).

Until now, social structure and reputation have rarely been considered together. Previous research has shown that either factor can help sustain cooperation even in short interactions. The valuable contribution of Buskens and Weesie (2000) is to how social structure and reputation can reinforce each other in sustaining cooperation. Information sharing that allows the formation of reputations allows cooperation to be sustained with short interaction in social structures that are less rigid than fixed geographic positions.

Rational and Adaptive Play

Hegselmann and Flache (2000) study the minimal conditions to sustain cooperation in an iterated Prisoner's Dilemma when the players are either rational or adaptive. With rational players, they consider the conditions needed to sustain cooperation if the players are using either Grim Trigger or Tit for Tat. As I discussed earlier, although Grim Trigger has less stringent requirements to sustain cooperation, I think a player in most situations would be ill-advised to use Grim Trigger because its lack of forgiveness can get it into a lot of trouble.

The innovative part of Hegselmann (2000) is its treatment of a particular kind of adaptive player. Building on the pioneering work of Bush and Mosteller (1955) and Rapoport and Chammah (1965) they define a specific learning strategy. This strategy changes its propensity to cooperate as a function of its own decision and the satisfaction it derived from the resulting outcome. Unlike most learning rules that have been studied

in the literature, they assume that an actor stops learning and becomes committed to a particular choice once its propensity to make that choice becomes sufficiently high. This can result in a mutual lock-in that provides the basis of some of their analytic results.

Unfortunately, the success of adaptive play is often highly dependent on the details of the learning rule itself, and especially on the strategies being used by the other agents in the population. This makes it hard to generalize about the value of adaptive approaches to playing the Prisoner's Dilemma. For example, in the two computer tournaments for the Prisoner's Dilemma that I ran (Axelrod 1984), there were a number of different learning rules submitted, some of them quite sophisticated. None of them did very well. They ran into two problems. First, the initial values of their propensities often implied that they would mix cooperation and defection until they gained substantial experience into the consequences of each. The defections in this mix of choices often got them into trouble with the other rules in the tournament. Second, the other rules often had trouble "making sense" of the probability mix of cooperation and defection used by the learning rules, and failing to make sense of it, they sometimes just gave up and defected for a while. The learning rule, in turn, was likely to draw the conclusion that the best thing to do if the other was defecting was to defect in turn, leading both sides to confirm their negative expectations of the other. In short, it is difficult to design a learning rule that will be effective with a wide range of other strategies, and not just with twins of oneself.

An alternative approach to designing a learning rule by hand is to let the entire population evolve based on survival of the most effective strategies. In effect, the players whose strategies are doing poorly learn from the players with more effective strategies.

There are two ways of doing this. The first is to use a fixed set of strategies and have them “reproduce” in proportion to their success in each “generation” (Axelrod 1984, pp. 50-52; Hofbauer and Sigmund 1998). The second is to allow the strategies to “reproduce” in proportion to their success, and allow new strategies to be introduced by means of mutation (Axelrod 1987, Lindgren 1991, Binmore and Samuelson 1992, Lomborg 1996). Letting the population of strategies evolve is generally a more robust way of studying adaptation than using a fixed set pre-specified strategies.

Envy

Lehno (2000) provides a defense of moderate envy. As he points out, one of my original pieces of advice to people who find themselves in an iterated Prisoner’s Dilemma was “don’t be envious” (Axelrod 1985, 110-3). To be envious, I meant to strive for a *greater* payoff than the other player.⁸ Lehno (2000) discusses two other meanings of envy. The first is a disposition to *avoid getting less* than the other player. The second meaning is more limited, and refers to the disposition to prevent others from doing better by unfair means.

I agree with Lehno that envy in the sense of demand for fairness is an important feature of human motivation. Indeed, one can make an evolutionary argument about why a strong disposition to insist on fairness might be part of our genetic heritage. After all, in highly competitive situations (such as allocation of scarce food or access to mates) letting

others get ahead could be detrimental to ones fitness (Buss 1999, 366f). The strong emotional drive to punish those who we envy might even have a fitness advantage by deterring exploitation, even if it is costly to us if evoked. As Frank (1988, p. 245) says, “The emotion of envy acts as a commitment device that prevents people from accepting profitable, but one-sided, transactions. Envious persons often behave irrationally, but there is genuine material advantage in being an effective bargainer.”⁹

Social sciences have a blind spot in regard to envy. For example, the massive *Handbook of Social Psychology* (Gilbert et al. 1998) has only a single mention of envy in 1,900 pages, and that sentence simply distinguishes envy from jealousy. This blind spot is nothing new. Schoeck (1966, p. 99) found that the first encyclopedic work on the behavioral sciences (Berelson and Steiner 1964) did not have a single index entry on envy. Why this blind spot? Certainly part of the answer is cultural. “In all cultures of mankind, in all proverbs and fairytales, the emotion of envy is condemned. The envious person is universally exhorted to be ashamed of himself.” (Schoeck 1966, p. 1) Yet there are many other emotions, which are condemned, and yet extensively studied, so it is not clear why envy is blind spot. Whatever the reason, Lehn, does a service by providing a strategic analysis of envy.

Lehno (2000) suggests that in the iterated Prisoner’s Dilemma a player should take care not to let the other side gain a one-sided advantage. To the extent that this means you should be provokable by a defection from the other player, I agree. I also

⁸ This is Dawkins’ (1989, p. 220) excellent formulation. I regret that I was not as clear as I should have been.

⁹ Recent experimental work on ultimatum games offers insight into the reluctance of people to accept unfair bargains. See for example Larrick R. P. and S. Blount (1997) and Huck and Oechssler (1999).

agree with Lehno that Tit for Tat's provocability can be considered to be equivalent to a moderate degree of envy (in his first sense) since it functions to prevent the other player from getting very far ahead.

Lehno (2000) considers strategies that deal with noise. It is well known that the Tit for Tat strategy suffers from even small amounts of noise because a single mistaken defection can echo indefinitely (Molander 1985). Three approaches have been proposed to deal with noise in the iterated Prisoner's Dilemma (see Wu and Axelrod 1995). The first two are variants of Tit for Tat. Generous Tit for Tat allows some percentage of the other player's defections to go unpunished. Contrite Tit for Tat avoids responding to the other player's defection after ones own unintended defection. A completely different approach is based on the learning principle that the same choice is repeated if and only if the most recent payoff was high (i.e., R or T). Wu and Axelrod (1995) show that Generous Tit for Tat and Contrite Tit for Tat both did well when noise was added to the environment of the second round of the Prisoner's Dilemma tournament.¹⁰ Thus when Tit For Tat is modified with generosity or contrition it remains a highly robust strategy in a noisy environment.

Lehno (2000) identifies a strategy he calls Moderate Envy. This is the strategy of defecting whenever the other player has defected more than oneself. In the absence of noise, this strategy is identical to Tit for Tat. In the presence of noise, it functions much like Contrite Tit for Tat by cooperating if it gains an "unearned" advantage. Lehno also identifies a strategy he calls Sophisticated Envy which is like Moderate Envy except that it tries to get out of a seemingly hopeless cycle of mutual defections by cooperating if the

other player gets quite far ahead. Whether Sophisticated Envy is robust in the sense of doing well with a wide variety of other strategies remains to be seen.

The heart of my previous advice about envy was the suggestion that comparing ones payoff with the payoff of the other player could easily become a self-defeating process. For example, if players tried to maximize the difference between their own score and the other player's score, they would be turning the game into a zero sum contest in which all opportunities to cooperate would vanish. A better standard of comparison is how well you are doing relative to how well someone else could be doing in your shoes (Axelrod 1984, p. 111). In my experience, people often fallaciously assume that the world is a zero-sum game like a sports contest. This is why it comes as such a surprise that a strategy such as Tit for Tat can win a tournament without doing better than anyone it meets (and not being envious in my sense). The primary value of studying and teaching the Prisoner's Dilemma is that it highlights the possibility that both sides can do well. One can interpret Tit for Tat as displaying moderate envy in Lehno's first sense (since it does not let the other player get very far ahead), but it is not envious in my sense (since it does not strive to do better than the player). Regardless of definitions, the key point is that the robust success of strategies that rely on reciprocity comes from their ability to elicit cooperation from a wide range of strategies.

Finally, the willingness to tolerate the success of others can be valuable for a society. As Frank (1999, p. 121) points out, "The explosive progress of the industrial economies of the West has been in no small measure the result of a generally shared cultural understanding that concerns about relative standing are simply not legitimate.

¹⁰ Neither the learning rule called Pavlov nor its generous variant did not do well in this

This is not to say that people in the capitalist societies never experience a twinge of envy or resentment when an acquaintance succeeds on a spectacular scale. It is just that such feelings have never been seen as a legitimate basis for restricting the options of others.”

Exit

The Edk-Group (2000) analyzes the effect of allowing players to exit from an unsatisfactory relationship. This extension of the standard Prisoner’s Dilemma is accomplished by including in the player’s strategy an option to end the bilateral relationship based on the history of the game so far. The analysis is conducted by computer simulation using a set of fifteen strategies specified by the authors. The population adapts over time by periodically having the least successful players give up their current strategy and adopt a randomly selected strategy from the specified set. The results show that in this setting there is the possibility of clever opportunism. Yet the most successful strategy is one that cooperates until the other player defects, and then immediately exits.

To assess the robustness of this result, they also arranged several settings in which a more limited set of specified strategies were used. Again, the opportunistic strategy often did fairly well, but the most successful strategy in most settings was the one that always cooperated and exited at the first defection by the other. Thus the possibility of exit tended to select against uncooperative players. As the Edk-Group (2000) point out, this conclusion is in line with the conclusions of slightly different simulations of exit by

environment (Wu and Axelrod 1995).

Schussler (1989 and 1990) and Majeski et al. (1997). On the other hand they also note that Ashlock et al. (1996) found that when there is preferential (rather than random) partner selection, cooperation is even more robust. If there is also a waiting penalty for exit, then the level of cooperation the population can sustain depends on size of the penalty for exit as well as social structure that determines partner selection (Macy and Skvoretz 1998). Together with other studies of the voluntary exit and ostracism (Hirshleifer and Rasmusen 1989, Epstein 1998, Stanley et al. 1995, Riolo 1997, Sherratt and Roberts 1998), the basis now exists to develop a deeper understanding of the role of mobility in sustaining or undermining cooperation.

Developing a deeper understanding of the consequences of a factor such as mobility will require that various studies be comparable in most ways, so that the effects of differences can be systematically assessed. Unfortunately, simulation studies allow researchers to vary so many details that it is often difficult to assess the causes of differences in their results. A helpful technique is to begin by replicating an earlier study, and only then adding something new to the model (Axelrod et al. 1996).

For example consider the consequences of two different adaptation rules. The adaptation rule of the Edk-Group (2000) has the lowest scoring 10% of players switch strategies, and when a player switches it adopts a randomly chosen strategy. The effect is that a “weak but safe” strategy that always scores a little below average will thrive. In contrast, a more widely used adaptation rule is the replicator dynamic which reproduces each strategy in proportion to its average score in the population (Axelrod 1984, 50-52; Hofbauer and Sigmund 1998). With the replicator dynamic, a “weak but safe” strategy will eventually die out, rather than thrive. Having selection pressure apply to all players

in proportion to their success is usually a more realistic assumption than having it apply only to the bottom decile.

Another problem of comparing simulation results is the somewhat arbitrary set of strategies selected for including in the population. Since the effectiveness of each strategy depends not only on its own characteristics, but also on the population of players it meets, the mix of strategies is important. The original computer tournaments (Axelrod 1984) dealt with this by generating a population of strategies each of which was designed by someone who wanted to win the tournament. Another method is to specify a large universe of potential strategies that can be specified in a certain language, start with an initial population drawn from this universe, and then let the population evolve by mutation as well as selection (Axelrod 1987, Lindgren 1991, Binmore and Samuelson 1992, Lomborg 1996). As pointed out earlier, letting the population of strategies evolve with mutation is generally a more robust way of studying adaptation than using a fixed set pre-specified strategies. Likewise, letting the population of strategies evolve with mutation would be a more robust way of studying the effects of exit.

Conclusion

The six papers in this symposium clearly demonstrate that Cooperation Theory continues to be a fruitful paradigm for the conduct of research on an ever-growing set of important theoretical questions. The symposium shows how using and extending the original paradigm of the two-person iterated Prisoner's Dilemma provides rich

possibilities for studying the effects a wide range of factors such as the timing of moves, hostage taking, social networks, adaptive play, envy, noise and mobility. In light of the extensive existing literature on related models dealing many of these factors, the time is now ripe for the comparison of results of closely related models on each factor, as well as for the continuing addition of new themes.

References

- Abell, Peter and Diane Reyniers (2000), "Generalized Reciprocity and Reputation in the Theory of Cooperation: A Framework," *Analyse and Kritik*, this issue.
- Allison and Zelikow(1999), *Essence of Decision: Explaining the Cuban Missile Crisis*, 2nd edition (New York: Longman).
- Ashlock, D. M. Schumacker, E. Stanley, L. Testfason (1996), Preferential Partner Selection in an Evolutionary Study of the Prisoner's Dilemma, *Biosystems* 37, 99-125.
- Axelrod, Robert (1981), "Emergence of Cooperation Among Egoists," *American Political Science Review* 75, 306-318.
- Axelrod, Robert (1984), *The Evolution of Cooperation* (NY: Basic Books).

Axelrod, Robert (1987), "The Evolution of Strategies in the Iterated Prisoner's Dilemma," in Lawrence Davis (ed.), *Genetic Algorithms and Simulated Annealing* (London: Pittman; Los Altos, Calif.: Morgan Kaufman), 32-41. Reprinted in Axelrod (1997a).

Axelrod, Robert (1997a), *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* (Princeton, NJ: Princeton University Press).

Axelrod, Robert (1997b), "Advancing the Art of Simulation in the Social Sciences," in Rosario Conte, Rainer Hegselmann and Pietro Terna (eds.), *Simulating Social Phenomena* (Berlin: Springer), 21-41.

Axelrod, Robert and Douglas Dion (1988), "The Further Evolution of Cooperation," *Science* 242 (9 December) 1385-1390.

Axelrod, Robert and William D. Hamilton (1981), "The Evolution of Cooperation," *Science* 211. (27 March), 1390-1396.

Axelrod, Robert, Robert Axtell, Joshua Epstein, and Michael D. Cohen (1996), Aligning Simulation Models: A Case Study and Results, *Computational and Mathematical Organization Theory* 1, 123-141.

Bendor, Jonathan and Piotr Swistak (1997), "The Evolutionary Stability of Cooperation," *American Political Science Review* 91, 290-307.

Berelson, Bernard and Gary A. Steiner (1964), *Human Behavior: An Inventory of Scientific Findings* (New York: Harcourt, Brace & World).

Binmore Kenneth G. and Larry Samuelson (1992), "Evolutionary Stability in Repeated Games Played by Finite Automata, *Journal of Economic Theory* 57, 278-305.

- Boerlijst, M. C., M. A. Nowak, K. Sigmund (1997), "The Logic of Contrition," *Journal of Theoretical Biology* 185, 281-293
- Bush, R. R. and F. Mosteller (1955), *Stochastic Models for Learning* (New York: Wiley).
- Buskens, Vincent and Jeroen Weesie (2000), "Cooperation via Social Networks," *Analyse and Kritik*, this issue.
- Buss, David M. (1999), *Evolutionary Psychology: The New Science of the Mind* (Boston: Allyn and Bacon).
- Cohen, Michael D., Rick Riolo and Robert Axelrod (1999), "The Emergence of Social Organization in the Prisoner's Dilemma: How Context-Preservation and other Factors Promote Cooperation," Santa Fe Institute Working Paper 99-01-002.
- Dawkins, Richard (1989), *The Selfish Gene*, new edition (Oxford and New York: Oxford University Press).
- Dewey, Horace (1988), "Russia's Debt to the Mongols in Suretyship and Collective Responsibility," *Comparative Studies in Society and History* 30, 249-270.
- Edk-Group (2000), "Exit, Anonymity and the Chances of Egotistical Cooperation," *Analyse and Kritik*, this issue.
- Epstein, Joshua M. (1998), Zones of Cooperation in Demographic Prisoner's Dilemma, *Complexity* 4, 36-48.
- Frank, Robert H. (1988), *Passions with Reason: The Strategic Role of Emotions* (New York: W. W. Norton,).
- Frean, Marcus R. (1994), "The Prisoner's Dilemma without Synchrony," *Proceedings of the Royal Society of London, Series B – Biological Sciences*, 257, no. 1348 (July 22), 75-77.

- Freidman, James W. (1971), "A Non-Cooperative Equilibrium in Supergames," *Review of Economic Studies* 38, 1-12.
- Fudenberg, Drew and David Levine, (1998), *Theory of Learning in Games* (Cambridge, Mass.: M.I.T. Press).
- Gilbert, Daniel T., Susan T. Fiske, and Gardner Lindzey (eds.) (1998), *The Handbook of Social Psychology*, 4th edition, two volumes (Boston, MA: McGraw-Hill).
- Grim, Patrick (1997), "The Greater Stability of the Spatialized Prisoner's Dilemma," *Journal of Theoretical Biology* 173, 353-359.
- Hamilton, William D. (1964), "The Genetical Evolution of Social Behavior (I and II)," *Journal of Theoretical Biology* 7, 1-16, 17-52.
- Hauert, Ch. and H. G. Schuster (1998), "Extending the Iterated Prisoner's Dilemma without Synchrony," *Journal of Theoretical Biology* 192, 155-166.
- Hegselmann, Rainer and Andreas Flache (2000), "Rational and Adaptive Playing: A Comparative Analysis for All Possible Prisoner's Dilemmas," *Analyse and Kritik*, this issue.
- Hirshleifer D, and E. Rasmusen (1989), Cooperation in a Repeated Prisoner's Dilemma with Ostracism, *Journal of Economic Behavior and Organization* 12, 87-106.
- Hofbauer, Josef and Karl Sigmund (1998), *Evolutionary Games and Population Dynamics* (Cambridge, UK: Cambridge University Press)
- Huck S. and J. Oechssler (1999), "The Indirect Evolutionary Approach to Explaining Fair Allocations," *Games and Economic Behavior* 28, 13-24.

- Larrick R. P. and S. Blount (1997), "The Claiming Effect: Why Players Are More Generous in Social Dilemmas Than in Ultimatum Games," *Journal of Personality and Social Psychology* 72, 810-825.
- Lehno, Bernd (2000), "In Defense of Moderate Envy," *Analyse und Kritik*, this issue.
- Leimar, Olaf (1997), "Repeated Games: A State Space Approach," *Journal of Theoretical Biology* 184, 471-498.
- Lindgren, K. and M. G. Nordahl (1994), "Evolutionary Dynamics of Spatial Games," *Physica D* 75, 292-309.
- Lindgren, Kristian (1991), "Evolutionary Phenomena in Simple Dynamics," in C. G. Langton et al. (eds.), *Artificial Life II* (Reading, MA: Addison-Wesley).
- Linster, Bruce G. (1992), "Evolutionary Stability in the Infinitely Repeated Prisoner's Dilemma Played by Two-State Moore Machines," *Southern Economic Journal* 58, 880-903.
- Lomborg, Bjorn, "Nucleus and Shield: The Evolution of Social Structure in the Iterated Prisoner's Dilemma," *American Sociological Review* 61, 278-307.
- Macy, M. and M. Skvoretz (1998), "The Evolution of Trust and Cooperation among Strangers," *American Sociological Review* 63, 638-660.
- Majeski, s., G. Linden, C. Linden, and A. Spitzer (1997), "A Spatialized Prisoner's Dilemma Game Simulation with Movement," in R. Conte, R. Hegselmann, and P. Terna (eds.), *Simulating Social Phenomena* (Berlin: Springer), 161-167.
- March, James G. (1978), "Bounded Rationality, Ambiguity and the Engineering of Choice," *The Bell Journal of Economics* 9, 596-608.

- Molander, Per (1985), "The Optimal Level of Generosity in a Selfish Uncertain Environment," *Journal of Conflict Resolution* 29, 611-618.
- Nakamura, M., H. Matusda and Y. Iwasa (1997), "The Evolution of Cooperation in a Lattice-Structured Population," *Journal of Theoretical Biology* 184, 65-81.
- Nowak, Martin A. and Karl Sigmund (1994), "The Alternating Prisoner's Dilemma," *Journal of Theoretical Biology* 168, 219-226.
- Nowak, Martin A. and Karl Sigmund (1998a), "The Dynamics of Indirect Reciprocity," *Journal of Theoretical Biology* 194, 561-574.
- Nowak, Martin A. and Karl Sigmund (1998b), "Evolution of Indirect Reciprocity by Image Scoring," *Nature* 393 (11 June), 573-577.
- Nowak, Martin A. and Karl Sigmund (1999), "Phage-Lift for Game Theory," *Nature* 398, 367-368.
- Nowak, Martin A. and Robert M. May (1992), "Evolutionary Games and Spatial Chaos," *Nature* 359 (29 October), 826-829.
- Nowak, Martin A., S. Bonhoeffer, and Robert M. May (1994), "Spatial Gamers and the Maintenance of Cooperation," *Proceeding of the National Academy of Sciences, USA*, 91, 4877-4881.
- Parker, C. (1977), "Reciprocal Altruism in Papio Anabis," *Nature* 265, 441-443.
- Pollock, Gregory B. (1989), "Evolutionary Stability of Reciprocity in a Viscous Lattice," *Social Networks* 11, 175-212.
- Powell, Robert (1999), *In the Shadow of Power: States and Strategies in International Politics* (Princeton, NJ: Princeton University Press).

- Rapoport, Anatol and Albert M. Chammah (1965), *Prisoner's Dilemma* (Ann Arbor, Mich.: University of Michigan Press).
- Raub, Werner and Jeroen Weesie (2000), "Cooperation via Hostages," *Analyse and Kritik*, this issue.
- Riolo, Rick L. (1997), "The Effects of Tag-Meditated Selection of Partners in Evolving Populations of the Iterated Prisoner's Dilemma," In Thomas Back (ed.), *Proceeding of the International Conference on Genetic Algorithms (ICGA-97)* (San Francisco: Morgan Kaufmann), 378-385.
- Samuelson, Larry (1997), *Evolutionary Games and Equilibrium Selection* (Cambridge, MA: M.I.T. Press).
- Schelling, Thomas C. (1963), *The Strategy of Conflict* (New York: Oxford University Press).
- Schelling, Thomas C. (1966), *Arms and Influence* (New Haven, Conn.: Yale University Press).
- Schelling, Thomas C. (1978), *Micromotives and Macrobehavior* (New York: W. W. Norton).
- Schoeck, Helmut (1966), *Envy: A Theory of Social Behavior* (New York: Harcourt Brace & World).
- Schussler, R. (1990), *Koopeation unter Egoisten* (Munich: Oldenbourg Verlag).
- Schussler, R. (1989), Exit Threats and Cooperation under Anonymity, *Journal of Conflict Resolution* 33, 727-749.
- Sherratt T. N. and G. Roberts (1998), The Evolution of Generosity and Choosiness in Cooperative Exchanges, *Journal of Theoretical Biology* 193, 167-177.

- Simon, Herbert A. (1982), *Models of Bounded Rationality*, vol. 1 and 2 (Cambridge, MA: M.I.T Press).
- Simon, Herbert A. (1990), "A Mechanism for Social Selection and Successful Altruism," *Science* 250 (21 December), 1665-1668.
- Stanley E. A., D. Ashlock, and M. Smucker (1995), Iterated Prisoner's Dilemma choice and Refusal of Partners: Evolutionary Results, *Lecture Notes in Artificial Intelligence* 929, 490-502.
- Trivers, Robert L. (1971), "The Evolution of Reciprocal Altruism," *Quarterly Review of Biology* 46, 35-57.
- Turner, P. E. and L. Chao, "Prisoner's Dilemma in an RNA Virus," *Nature* 398, 441-443.
- Wedekind, Claus and Manfred Milinski (1996), "Human Cooperation in the Simultaneous and Alternating Prisoner's Dilemma: Pavlov versus Generous Tit-for-Tat," *Proceedings of the National Academy of Sciences, USA* 93, 2686-2689.
- Wu, Jianzhong and Robert Axelrod (1995), "How To Cope With Noise in the Iterated Prisoner's Dilemma," *Journal of Conflict Resolution* 39, 183-189.
- Young, H. Peyton (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton NJ: Princeton University Press).