ROBERT AXELROD

The Complexity of Cooperation

Agent-Based Models of Competition and Collaboration

PRINCETON STUDIES IN COMPLEXITY

Introduction

THE TITLE of this book illustrates the dual purposes of the volume. One meaning of "The Complexity of Cooperation" refers to the addition of complexity to the most common framework for studying cooperation, namely the two-person iterated Prisoner's Dilemma. Adding complexity to that framework allows the exploration of many interesting and important features of competition and collaboration that are beyond the reach of the Prisoner's Dilemma paradigm.

The second meaning of "The Complexity of Cooperation" refers to the use of concepts and techniques that have come to be called complexity theory. Complexity theory involves the study of many actors and their interactions. The actors may be atoms, fish, people, organizations, or nations. Their interactions may consist of attraction, combat, mating, communication, trade, partnership, or rivalry. Because the study of large numbers of actors with changing patterns of interactions often gets too difficult for a mathematical solution, a primary research tool of complexity theory is computer simulation. The trick is to specify how the agents interact, and then observe properties that occur at the level of the whole society. For example, with given rules about actors and their interactions, do the actors tend to align into two competing groups? Do particular strategies dominate the population? Do clear patterns of behavior develop?

The simulation of agents and their interactions is known by several names, including agent-based modeling, bottom-up modeling, and artificial social systems. Whatever name is used, the purpose of agent-based modeling is to understand properties of complex social systems through the analysis of simulations. This method of doing science can be contrasted with the two standard methods of induction and deduction. Induction is the discovery of patterns in empirical data.[1] For example, in the social sciences induction is widely used in the analysis of opinion surveys and macroeconomic data. Deduction, on the other hand, involves specifying a set of axioms and proving consequences that can be derived from those assumptions. The discovery of equilibrium results in game theory using rational-choice axioms is a good example of deduction.

statistics

science

Agent-based modeling is a third way of doing science. Like deduction, it starts with a set of explicit assumptions. But unlike deduction, it does

[1] Induction as a search for patterns in data should not be confused with mathematical induction, which is a technique for proving theorems.

not prove theorems. Instead, an agent-based model generates simulated data that can be analyzed inductively. Unlike typical induction, however, the simulated data come from a rigorously specified set of rules rather than direct measurement of the real world. Whereas the purpose of induction is to find patterns in data and that of deduction is to find consequences of assumptions, the purpose of agent-based modeling is to aid intuition.

Agent-based modeling is a way of doing thought experiments. Although the assumptions may be simple, the consequences may not be at all obvious. Numerous examples appear throughout this volume of locally interacting agents producing large-scale effects. The large-scale effects of locally interacting agents are called "emergent properties" of the system. Emergent properties are often surprising because it can be hard to anticipate the full consequences of even simple forms of interaction.[2]

There are some models, however, in which emergent properties can be formally deduced. Good examples include the neoclassical economic models in which rational agents operating under powerful assumptions about the availability of information and the capability of optimizing can achieve an efficient reallocation of resources among themselves through costless trading. But when the agents use adaptive rather than optimizing strategies, deducing the consequences is often impossible; simulation becomes necessary.

Throughout the social sciences today, the dominant form of modeling is based upon the rational-choice paradigm. Game theory, in particular, is typically based upon the assumption of rational choice. In my view, the reason for the dominance of the rational-choice approach is not that scholars think it is realistic. Nor is game theory used solely because it offers good advice to a decision maker, because its unrealistic assumptions undermine much of its value as a basis for advice. The real advantage of the rational-choice assumption is that it often allows deduction.

The main alternative to the assumption of rational choice is some form of adaptive behavior. The adaptation may be at the individual level through learning, or it may be at the population level through differential survival and reproduction of the more successful individuals. Either way, the consequences of adaptive processes are often very hard to deduce when there are many interacting agents following rules that have nonlinear effects. Thus the simulation of an agent-based model is often the only viable way to study populations of agents who are adaptive rather than fully rational.

Although agent-based modeling employs simulation, it does not aim to

[2] Some complexity theorists consider surprise to be part of the definition of emergence, but this raises the question of surprising to whom?

provide an accurate representation of a particular empirical application. Instead, the goal of agent-based modeling is to enrich our understanding of fundamental processes that may appear in a variety of applications. This requires adhering to the KISS principle, which stands for the army slogan "keep it simple, stupid."

The KISS principle is vital because of the character of the research community. Both the researcher and the audience have limited cognitive ability. When a surprising result occurs, it is very helpful to be confident that we can understand everything that went into the model. Although the topic being investigated may be complicated, the assumptions underlying the agent-based model should be simple. The complexity of agent-based modeling should be in the simulated results, not in the assumptions of the model.

Of course there are many other uses of computer simulation in which the faithful reproduction of a particular setting is important. A simulation of the economy aimed at predicting interest rates three months into the future needs to be as accurate as possible. For this purpose the assumptions that go into the model may need to be quite complicated. Likewise, if a simulation is used to train the crew of a supertanker or to develop tactics for a new fighter aircraft, accuracy is important and simplicity of the model is not. But if the goal is to deepen our understanding of some fundamental process, then simplicity of the assumptions is important, and realistic representation of all the details of a particular setting is not.

My earlier work on the Prisoner's Dilemma (Axelrod 1984) illustrates this theme. My main motivation for learning about effective strategies was to find out how cooperation could be promoted in international politics, especially between the East and the West during the Cold War. As it happened, my tournament approach and the evolutionary analysis that grew out of it suggested applications of which I had not even dreamed. For example, controlled experiments show that stickleback fish use the TIT FOR TAT strategy to achieve cooperation based upon reciprocity (Milinski 1987).

At a political science convention, a colleague came up to me and said she really appreciated my work and found it helpful for her divorce. She explained that my book showed her that she had been a sucker during her marriage, always giving in to her husband. I asked whether the book helped save her marriage. "No," she replied. "I didn't want to save my marriage. But it certainly helped with the divorce settlement. I started to play TIT FOR TAT, and once he learned that I couldn't be pushed around, I got a much better deal."

The fact that a single model, in this case the Prisoner's Dilemma, can be useful in understanding the dynamics between foraging fish and between divorcing people is not due to the accuracy of the model in representing the details of either situation. Instead it is due to the fact that an extremely simple model captures a fundamental feature of many interactions. What the Prisoner's Dilemma captures so well is the tension between the advantages of selfishness in the short run versus the need to elicit cooperation from the other player to be successful in the longer run. The very simplicity of the Prisoner's Dilemma is highly valuable in helping us to discover and appreciate the deep consequences of the fundamental processes involved in dealing with this tension.

A moral of the story is that models that aim to explore fundamental processes should be judged by their fruitfulness, not by their accuracy. For this purpose, realistic representation of many details is unnecessary and even counterproductive. The models presented in the volume follow this same logic of simplicity. The intention is to explore fundamental social processes. Although a particular application may have motivated a given model, the primary aim is to undertake the exploration in a manner so general that many possible settings could be illuminated.

Taken as a whole, this book presents a set of studies that are unified in three ways. 1) First, they all deal with problems and opportunities of cooperation in a more or less competitive environment. Second, they all employ 2) models that use adaptive rather than rational agents. Although people may try to be rational, they can rarely meet the requirements of information or foresight that rational models impose (Simon 1955; March 1978). Third, they all use computer simulation to study the emergent 3) properties of the interactions among the agents. Thus they are all agent-based models. The simulation is necessary because the interactions of adaptive agents typically lead to nonlinear effects that are not amenable to the deductive tools of formal mathematics.

The chapters can be read either separately or as a whole. The order of the presentation represents a progression from variations on the Prisoner's Dilemma paradigm (Chapters 1 and 2), to different strategic models (Chapters 3, 4, and 5), to an examination of the emergence of new political actors and shared culture (Chapters 6 and 7). The order of the chapters is also the order in which I did the work, with the exception that Chapter 2 represents later work on an earlier theme.

The first project represents my effort to go beyond the tournament approach to generating a rich strategic environment. The tournament approach solicited entries from professionals and amateurs, each trying to develop a strategy for the Prisoner's Dilemma that would do well in the environment provided by all the submissions. Having done two rounds of the tournament, I wondered whether the amount of cooperation I observed was due to the prior expectations of the people who submitted the rules. Fortunately, a colleague, John Holland, had developed

an automated method for evolving a population of strategies from a random start. The technique is called the genetic algorithm. I tried it, and it performed far beyond my expectations. The results are in Chapter 1.

An important extension of the basic Prisoner's Dilemma is consideration of what happens when a player might misunderstand what the other did on the previous move or might fail to implement the intended choice. These kinds of "noise" can have a big impact on the performance of a given strategy, and hence on the best means of attaining cooperation among egoists. Several suggestions had been proposed in the literature for dealing with noise, including adding generosity or contrition to reciprocity, as well as a completely different strategy based upon learning through reward and punishment. I wanted to see how these different approaches would work in a noisy environment. A postdoctoral visitor from China, Wu Jianzhong, and I found that generous or contrite versions of the classic TIT FOR TAT strategy did very well in a variegated noisy environment, even better than the Pavlovian strategy. Chapter 2 explains how these strategies performed and why.

For a long time, I had been eager to move beyond the two-person format of the basic Prisoner's Dilemma. I especially wanted to find out how cooperation could emerge when many people interacted with each other in groups rather than in pairs. It was well known that the straightforward extension of the Prisoner's Dilemma to an *n*-person version will not sustain cooperation very well because the players have no way of focusing their punishment on someone in the group who has failed to cooperate. Nevertheless, social norms do emerge and are often quite powerful means of sustaining cooperation. So I developed a "norms game" that allowed players to punish individuals who do not cooperate. It turned out that another twist was needed lest all the cooperators be tempted to let someone else be the one to bear the costs of disciplining the noncooperators. This led to a wide-ranging study of the mechanisms for promoting norms (Chapter 3).

Another form of cooperation occurs when people organize themselves into groups to compete with each other. This is clearly an example of collaboration in the interests of competition. It takes place in many forms, including alliances among nations, strategic partnerships among businesses, and coalitions among political parties in parliamentary democracies. Having worked on the problem of coalition formation in Italy as part of my dissertation in the late 1960's, I was struck by how political parties wanted to work with others who were similar to themselves (Axelrod 1970). Two decades later, I returned to this theme of choosing sides based upon affinity rather than strategic advantage. Working with a graduate student, Scott Bennett, I developed a model for how players choose sides. We found that the model actually did a good job of accounting for how European countries were aligned in World War II (Chapter 4).

The same model even worked well in accounting for how computer companies took sides in the competition to develop standards for the UNIX operating system (Chapter 5). This was work done with Scott Bennett and three collaborators from the Michigan Business School: Will Mitchell, Robert E. Thomas, and Erhard Bruderer.

An even deeper problem is how independent actors sometimes cooperate to such an extent that they give up most of their independence. The result is a new level of organization that behaves as an independent actor in its own right. Multicellular organisms evolved this way, and so have many large business organizations. My approach to analyzing how new levels of political actors can arise uses a model of war, threats, and commitments. The agent-based model and its results are provided in Chapter 6.

Whereas the model in Chapter 6 attributes the emergence of new actors to the dynamics of coping with conflict, I also wanted to study an even more fundamental question: how people become more alike so that they find it easier to work together in the first place. This led to a study of the process of social influence and the emergence of shared culture. Once again, the transformations of the post-Cold War environment helped emphasize the importance of returning to some very fundamental issues about the basis for cooperation within as well as between societies. The resulting model of social influence and cultural change is given in Chapter 7.

Two appendixes provide supporting material about agent-based modeling. Appendix A develops the concepts and methods of comparing agent-based models through a process called "alignment." Alignment is needed to determine whether two modeling systems can produce the same results, which in turn is the basis for critical experiments and for tests of whether one model can subsume another. The work provides a case study of alignment, using the model of social influence presented in Chapter 7. The project was done with Robert Axtell, Joshua Epstein, and Michael Cohen. Appendix B provides resources for students and scholars who wish to do their own agent-based modeling. It includes advice on programming such models, exercises to develop one's skills, and suggested readings for applications of complexity theory and agent-based modeling to the social sciences.

Associated with this volume is an Internet site.[3] The site includes the source code and documentation for most of the models in this book. It also provides links to many topics related to complexity theory, agent-based modeling, and cooperation.

[3] http://pscs.physics.lsa.umich.edu/Software/ComplexCoop.html

!!!

References

Axelrod, Robert. 1970. *Conflict of Interest.* Chicago: Markham.

______. 1984. *The Evolution of Cooperation.* New York: Basic Books.

March, James G. 1978. "Bounded Rationality, Ambiguity and the Engineering of Choice." *Bell Journal of Economics* 9: 587–608.

Milinski, Manfred. 1987. "TIT FOR TAT in Sticklebacks and the Evolution of Cooperation." *Nature* 23: 434–35.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69: 99–118.