Advancing the Art of Simulation in the Social Sciences

Obtaining, analyzing, and sharing results of computer models

BY ROBERT AXELROD

Robert Axelrod is a political scientist best known for his book The Evolution of Cooperation. *His most recent book is* The Complexity of Cooperation *(Princeton University Press, 1997). He received a MacArthur Prize Fellowship and was elected to the National Academy of Sciences. He teaches at the University of Michigan, where he is the Arthur Bromage Distinguished University Professor of Political Science and Public Policy.*

INTRODUCTION

Simulation is a young and rapidly growing field in the social sciences. As in most young fields, its promise is greater than its accomplishments. The purpose of this article is to suggest what is needed for the field to become mature so that its potential contribution to the social sciences can be realized.

The next section discusses the variety of purposes that simulation can serve, emphasizing the discovery of new principles and relationships. Advice is then offered for how to perform research with simulation. Topics include programming a simulation model, analyzing the results, and sharing the results. Next, the frequently neglected topic of replication is considered, with detailed descriptions of two replication projects. The final section suggests how to advance the art of simulation by building a community of social scientists and others who use computer simulation in their research.

THE VALUE OF SIMULATION

Let us begin with a definition of simulation. "Simulation means driving a model of a system with suitable inputs and observing the corresponding outputs" [1]. While this definition is useful, it does not suggest the diverse purposes for which simulation can be used. These purposes include prediction, performance, training, entertainment, education, proof, and discovery.

1. *Prediction.* Simulation is able to take complicated inputs, process them by accounting for hypothesized mechanisms, and then generate their consequences as predictions. For example, if the goal is to predict interest rates in the economy three months into the future, simulation can be the best available technique.

2. *Performance.* Simulation can also be used to perform certain tasks. This is typically the domain of artificial intelligence. Tasks to be performed include medical diagnosis, speech recognition, and function optimization. To the extent that artificial intelligence techniques mimic the way humans deal with these tasks, the artificial intelligence method can be thought of as simulation of human perception, decision making, or social interaction. To the extent that artificial intelligence techniques exploit the special strengths of digital computers, simulations of task environments can also help design new techniques.

3. *Training.* Many of the earliest and most successful simulation systems were designed to train people by providing a reasonably accurate and dynamic interactive representation of a given environment. Flight simulators for pilots is an important example of the use of simulation for training.

CCC 1076-2787/97/02016-07

Advancing the state of the art of simulation in the social sciences requires appreciating the unique value of simulation as a research methodology. This article offers advice on simulation research, focusing on the programming of a simulation model, analyzing the results, and *sharing* the results. Replicating other simulations is emphasized, with examples of the procedures and difficulties involved in the process of replication. Finally, suggestions are offered for building a community of social scientists who use simulation.

4. Entertainment. From training, it is only a small step to entertainment. Flight simulations on personal computers are fun; so are simulations of completely imaginary worlds.

5. Education. From training and entertainment, it is only a small step to the use of simulation for education. A good example is the computer game SimCity, an interactive simulation allowing the user to experiment with a hypothetical city by changing many variables, such as tax rates and zoning policy. For educational purposes, a simulation need not be rich enough to suggest a completely real or imaginary world. The primary use of simulation in education is to allow users to learn relationships and principles for themselves.

6. Proof. Simulation can be used to provide an existence proof. For example, Conway's Game of Life demonstrates that extremely complex behavior can result from very simple rules [2].

7. Discovery. As a scientific methodology, simulation's value lies principally in prediction, proof, and discovery. Using simulation for prediction can help validate or improve the model upon which the simulation is based. Prediction is the use that most people think of when they consider simulation as a scientific technique. But the use of simulation for the discovery of new relationships and principles is as important as proof or prediction. In the social sciences in particular, even highly complicated simulation models can rarely prove to be completely accurate. Physicists have accurate simulations of the motion of electrons and planets, but social scientists are not as successful in accurately simulating the movement of workers or armies. Nevertheless, social scientists have been successful in using simulation to discover important relationships and principles from very simple models. Indeed, as discussed below, the simpler the model, the easier it may be to discover and understand the subtle effects of its hypothesized mechanisms.

Schelling's simulation of residential tipping provides a good example of a simple model that provides important insight into a general process [3,4]. The model assumes that a family will move only if more than one third of its immediate neighbors are of a different type (e.g., race or ethnicity). The result is that very segregated neighborhoods form even though everyone is initially placed at random and everyone is somewhat tolerant.

To appreciate the value of simulation as a research methodology, one should think of it as a new way of conducting scientific research. Simulation can be compared with the two standard methods of induction and deduction. Induction is the discovery of patterns in empirical data. For example, in the social sciences, induction is widely used in the analysis of opinion surveys and the 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.

Simulation is a third research methodology. Like deduction, it starts with a set of explicit assumptions. But unlike deduction, it does not prove theorems. Instead, simulation generates data that can be analyzed inductively. Unlike typical induction, however, the simulated data come from a specified set of rules rather than direct measurement of the real world. While induction can be used to find patterns in data, and deduction can be used to find consequences of assumptions, simulation modeling can be used to aid intuition.

Simulation is a way of performing thought experiments. While the assumptions may be simple, the consequences may not be at all obvious. The large-scale effects of locally interacting agents are called "emergent properties" of the system. Emergent properties are often surprising because it can be difficult to anticipate all the consequences of even simple forms of interaction. (Some complexity theorists consider surprise to be part of the definition of emergence, but this raises the question, surprising to whom?)

There are some models, however, in which emergent properties can be formally deduced. Examples include the neo-classical economic models in which rational agents operating under powerful assumptions about the availability of information and the capability to optimize 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 on the rational choice paradigm. Game theory, in particular, is typically based on the assumption of rational choice. 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, since 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 for 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 difficult to deduce when there are many interacting agents following rules that have nonlinear effects. Thus, simulation is often the only viable way to study populations of agents who are adaptive rather than fully rational. While people may try to be rational, they can rarely meet the requirement of information or foresight that rational models impose [5,6]. One of the main advantages of simulation is that it allows the analysis of adaptive as well as rational agents.

An important type of simulation in the social sciences is agent-based modeling. This type of simulation is characterized by the existence of many agents who interact with each other with little or no central direction. The emergent properties of an agent-based model are then the result of "bottom-up" processes, rather than "top-down" direction.

An important type of simulation in the social sciences is agent-based modeling. This type of simulation is characterized by the existence of many agents who intereact with each other with little or no central direction.

Although agent-based modeling employs simulation, it does not necessarily aim to 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. It is important to keep the model as simple as possible. When a surprising result occurs, it is very helpful to be confident that one can understand everything that contributed to the model. Simplicity is also helpful in giving other researchers a realistic chance of extending one's model in new directions. While 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.

As pointed out earlier, there are 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.

SIMULATION RESEARCH

In order to advance the art of simulation in the social sciences, it is necessary to do more than consider the purpose of simulation. It is also necessary to be more self-conscious about the process of performing the research itself. To do so requires looking at three specific aspects of the research process which take place once the conceptual model is developed: the programming of the model, the analysis of the data, and the sharing of the results.

Programming a Simulation Model

The first question people usually ask about programming a simulation model is, "What language should I use?" My recommendation is to use one of the modern procedural languages, such as Pascal, C,or C++. For small projects, it may be easiest to program within a graphics or statistical package, or even a spreadsheet.

The programming of a simulation model should achieve three goals: validity, usability, and extendibility.

The goal of validity is for the program to correctly implement the model. This kind of validity is called "internal validity." Whether the model itself is an accurate representation of the real world is another kind of validity that is not considered here. Achieving internal validity is more difficult than it might seem. The problem is knowing whether an unexpected result is a reflection of a mistake in the programming or a surprising consequence of the model itself. For example, in one of my own models, a result was so counterintuitive that careful analysis was required to confirm that the result was a consequence of the model, and not due to a bug in the program. As is often the case, confirming that the model was correctly programmed was substantially more work than programming the model in the first place.

The goal of usability is to allow the researcher and those who run the program to interpret its output and understand how it works. Modeling typically generates a series of programs, each version differing from the others in a variety of ways. Versions can differ, for example, in which data are produced, which parameters are adjustable, and which rules govern agent behavior. Keeping track of all this is not trivial, especially when one compares new results with output of an earlier version of the program to determine exactly what might account for the differences.

The goal of extendibility is to allow a future user to adapt the program for new uses. For example, after writing a paper using the model, the researcher might want to respond to a question about what would happen if a new feature were added. In addition, another researcher might want to modify the program to test a new variant of the model. A program is much more likely to be extendible if it is written and documented with this goal in mind.

Analyzing the Results

Simulation typically generates huge amounts of data. In fact, one of the advantages of simulation is that if there is not enough data, one can always run the simulation again. Moreover, there are no problems of missing data or uncontrolled variables as there are in experimental or observational studies.

Despite the purity and clarity of simulation data, the analysis poses real challenges. Multiple runs of the same model can differ from each other due to differences in initial conditions and stochastic events. A major challenge is that results are often path-dependent, meaning that history matters. To understand the results often means understanding the details of the history of a given run. There are at least three ways in which history can be described.

First, history can be told as “news,” following a chronological order. For example, a simulation of international politics might describe the sequence of key events such as alliances and wars. This is the most straightforward type of story telling but often offers little in explanatory power.

Second, history can be told from the point of view of a single actor. For example, one could select just one of the actors and do the equivalent of telling the story of the rise and fall of the Roman Empire. This is often the easiest kind of history to understand and can be very revealing about the ways in which the model’s mechanisms have their effects over time.

Third, history can be told from a global point of view. For example, one would describe the distribution of wealth over time to analyze the extent of inequality among the agents. Although the global point of view is often the best for seeing large-scale patterns, the more detailed histories are often needed to determine the explanation for these large patterns.

While the description of data as history is important for discovering and explaining patterns in a particular simulation run, the analysis of simulations all too often stops there. Since virtually all social science simulations include some random elements in their initial conditions and in the operation of their mechanisms for change, the analysis of a single run can be misleading. In order to determine whether the conclusions from a given run are typical, it is necessary to do several dozen simulation runs using identical parameters (with different random number seeds) to determine which results are typical and which are unusual. While it may be sufficient to describe detailed history from a single run, it is also necessary to perform a statistical analysis of a set of runs to determine whether the inferences being drawn from the illustrative history are really well-founded. The ability to do this is yet one more advantage of simulation: The researcher can rerun history to see whether particular patterns observed in a single history are idiosyncratic or typical.

By using simulation, one can do even more than compare multiple histories generated from identical parameters. One can also systematically study the effects of changing the parameters. For example, agents can be given either equal or unequal initial endowments to see what impact this has over time. Likewise, the differences in mechanisms can be studied by performing systematic comparisons of different versions of the model. For example, in one version. agents might interact at random whereas in another version the agents might be selective in who they interact with. As in the simple change in parameters, the effects of changes in the mechanisms can be assessed by running controlled experiments with sets of simulation runs. Typically, the statistical method for studying the effects of these changes will be regression if the changes are quantitative, and analysis of variance if the changes are qualitative. As always in statistical analysis, two questions need to be distinguished and addressed separately: Are the differences statistically significant (i.e., not likely to have been caused by chance), and are the differences substantively significant (i.e., large enough in magnitude to be important)?

...one more advantage of simulation: The researcher can rerun history to see whether particular patterns observed in a single history are idiosyncratic or typical.

Sharing the Results

After cycling through several iterations of constructing the model, programming the simulation, and performing the data analysis, the final step in the research is sharing the results. As in most fields of research, the primary method of sharing results is through publication, most often in refereed journals or chapter-length reports in edited collections. In the case of social science simulation, there are several limitations with relying on this mode of sharing information. The basic problem is that it is hard to present a social science simulation briefly. There are at least three reasons.

First, simulation results are typically quite sensitive to the details of the model. Therefore, unless the model is described in great detail, the reader is unable to replicate or even fully understand what was done. Articles and chapters are often not long enough to present the full details of the model. (The issue of replication is addressed at greater length below.)

Second, the an-alysis of the results often includes some narrative description of histories of one or more runs, and such narrative is often lengthy. While statistical analysis can usually be described quite briefly in numbers, tables, or figures, the presentation of how inferences are drawn from the study of particular histories usually cannot be brief. This is due primarily to the amount of detail required to explain how the model’s mechanisms played out in a particular historical context. In addition, the paucity of well-known concepts and techniques for the presentation of historical data in context means that the writer cannot communicate this kind of information very efficiently. Compare this lack of shared concepts with the mature field of hypothesis testing in statistics. The simple phrase “$p < .05$” stands for the sentence, “The probability that this result (or a more extreme result) would have happened by chance is less than 5%.” Perhaps over time, the community of social science modelers will develop a collection of standard concepts that can become common knowledge and then be communicated briefly, but this is not true yet.

Third, simulation results often address an interdisciplinary audience. When this is the case, the unspoken assumptions and short-hand terminology that provide shortcuts for every discipline may need to be explicated at length to explain the motivation and premises of the work to a wider audience.

Fourth, even if the audience is of a single discipline, the computer simulations are still new enough in the social sciences that it may be necessary to explain very carefully both the power and the limitations of the methodology each time a simulation report is published.

Since it is difficult to provide a complete description of a simulation model in an article-length report, other forms of sharing information have to be developed. Complete documentation would include the source code for running the model, a full description of the model, how to run the program, and how to understand the output files. An established way of sharing this documentation is to mail a hard copy or a disk to anyone who asks the author for it. Another way is to place the material in an archive, such as the Interuniversity Consortium for Political and Social Research at the University of Michigan. This is already common practice for large empirical data sets such as public opinion surveys. Journal publishers could also maintain archives of material supporting their own articles. The archive then handles the distribution of materials, perhaps for a fee.

Two new methods of distribution are available: CD-ROM and the Internet. Each has its own characteristics worth considering.

A CD-ROM is suitable when the material is too extensive to distribute by traditional means or when it would be too time-consuming for a user to download it from the World Wide Web. A good example is animations of multiple simulation runs. The primary disadvantage is the cost to the user of purchasing the CD-ROM, either as part of the price of a book or as a separate purchase from the publisher.

Another new method is to place the documentation on the Internet. Today, the Web provides the most convenient way to use the Internet. By using the Internet for documentation, the original article need only provide the address of the site where the material is kept. This method has many advantages.

Unlike paper printouts, the material is available in machine-readable form. Using the Web makes the material immediately available from virtually anywhere in the world, with little or no effort required to answer each new request. Further, material on the Web can be readily updated and can be structured with hyperlinks to make clear the relationship between the parts. Material on the Web can also be easily cross-referenced from other Web sites. This is especially helpful since, as noted earlier, social science simulation articles are published in a wide variety of journals. As specialized Web sites develop to keep track of social science simulations, they can become valuable tools for the student or researcher who wants to find out what is available (e.g., www.econ.iastate.edu/tesfatsi/abe.htm, maintained by Leigh Tesfatsion, Iowa State University, and specializing in agent-based computational economics with pointers to simulation work in other fields).

A significant problem with placing documentation on the Web is how to guarantee it will still be there years later. Web sites tend to have high turnover. Yet a reader who comes across a simulation article ten years after publication should still be able to access the documentation. There are no well-established methods of guaranteeing that a particular Web server (e.g., at a university department) will maintain a given set of files for a decade or more. Computer personnel come and go, equipment is replaced, and budgetary priorities change. The researcher needs to keep a copy in case something happens to the Web server being used.

...placing documentation and perhaps even executable programs on the Internet has so many advantages that it is likely to become an important means of providing material needed to supplement the publication of simulation research.

The Internet offers more than just a means of documenting a simulation. It also offers the ability for a user to run a simulation program on his or her own computer. This can be done through a programming environment such as Java which allows the code that resides on the author's machine to be executed on the user's machine. A major advantage of this method of distributing a simulation program is that the same code can be run on virtually any type of computer. A good example is a simulation of a model of the spread of HIV infection. The description of the model, an article about its motivation, and a working version that can be run and even adapted by a distant user are all available on the Web. (The site is http://www.nytimes.com/library/cyber/week/1009aids.html.)

One disadvantage of using Java is that it is slower in execution than a locally compiled program. Another disadvantage of using Java or a similar programming environment is that there is no guarantee that the standards will be stable enough to allow easy use in ten years.

Despite the need to assure the durability of one's own Web site, placing documentation and perhaps even executable programs on the Internet has so many advantages that it is likely to become an important means of providing material needed to supplement the publication of simulation research. (Documentation and source code for many of my own agent-based models are on the Web at http://pscs.physics.lsa.umich.edu/Software/ComplexCoop. html.)

REPLICATION OF SIMULATIONS

Three important stages of the research process for performing simulation in the social sciences have been considered so far: programming, analyzing, and sharing computer simulations. All three are done for virtually all published simulation models. There is, however, another stage of the research process that is rarely done, but which needs to be considered: replication. New simulations are produced all the time, but rarely does anyone stop to replicate the results of anyone else's simulation model.

Replication is one of the hallmarks of cumulative science. It is needed to confirm whether the claimed results of a given simulation are reliable in the sense that they can be reproduced by someone starting from scratch. Without this confirmation, it is possible that some published results are simply incorrect due to programming errors, misrepresentation of what was actually simulated, or errors in analyzing or reporting the results. Replication can also be useful for testing the robustness of inferences from models. Finally, replication is needed to determine whether one model can subsume another, in the sense that Einstein's treatment of gravity subsumes Newton's.

Because replication is rare, it may be helpful to describe the procedures and lessons from two replication projects that I have been involved with. The first reimplemented one of my own models in a different simulation environment. The second sought to replicate a set of eight diverse models using a common simulation system.

The first replication project grew out of a challenge posed by Michael Cohen: Could a simulation model written for one purpose be aligned or "docked" with a general-purpose simulation system written for a different purpose. We chose my own cultural change model [7] as the target model for replication. For the general purpose simulation system, we chose the Sugarscape system developed by Joshua Epstein and Rob Axtell [8]. We invited Epstein and Axtell to modify their simulation system to replicate the results of my model. Along the way, the four of us discovered a number of interesting lessons, which are described in [9].

This systematic replication study demonstrated that replication is a feasible, although rarely performed, part of the process of advancing computer simulation in the social sciences. The lessons suggest that further replication would be worthwhile. The concepts and methods developed for this particular study suggest how further replications could be performed. Our finding that seemingly small differences were significant suggests that it is worth finding out whether this experience was typical. In particular, it would pay to replicate a diverse set of simulation models to see what types of problems arise.

Michael Cohen, Rick Riolo, and I took up this challenge. We selected a set of eight core models to replicate. We selected these models using six criteria: 1) simplicity (for ease of implementation, explanation, and understanding); 2) relevance to the social sciences; 3) diversity across disciplines and types of models; 4) reasonably short run times; 5) established heuristic value; and 6) accessibility through published accounts. Most of the eight models met at least five of these criteria.

Cohen, Riolo, and I implemented each of the models in the Swarm simulation system developed at Santa Fe Institute under the direction of Chris Langton. In each case, we identified the key results from the original simulations and determined what comparisons would be needed to test for equivalence. After a good deal more work than we had expected would be necessary, we were able to attain relational equivalence on all eight models. In most cases, the results were so close that we probably attained distributional equivalence as well, although we did not perform the statistical tests to confirm this.

Who should be better able to build new institutions than the researchers who use simulation to study real and potential societies?

We hoped to find some building blocks that were shared by several of these models that could provide the basis for a set of useful simulation techniques. Instead, we found little overlap. On the other hand, Riolo and Ted Belding developed a useful tool for running batch jobs of a simulation program to execute experimental designs. (This tool, called Drone, automatically runs batch jobs of a simulation program in Unix. It sweeps over arbitrary sets of parameters, as well as multiple runs for each parameter set, with a separate random seed for each run. The runs may be executed either on a single computer or over the Internet on a set of remote hosts. See http://pscs.physics.lsa.umich.edu/software/drone/index.html.)

The most important discovery we made in replicating the models is just how many things can go wrong.

The list below does not include the errors that we made in reimplementing the models, since the discovery and elimination of our own errors are just part of the normal process of debugging programs before they are regarded as complete and ready for publication. Instead, the list includes the problems we found in the published accounts or the programs that they describe. It should be noted that while these problems made it more difficult for us to replicate the original results, in no case did they make a major difference in the conclusions of the published accounts.

The first category of replication problem was ambiguity in the published descriptions. Ambiguities occurred in the description of the model and in the presentation of the numerical results. Ambiguities in the description of the model included the order in which the agents should be updated and what to do when there was a tie. Ambiguities in the description of the model included the meaning of a variable in a figure and the divisor used in a table. Some of these ambiguities in the published descriptions were resolved by seeing which of two plausible interpretations reproduced the original data. This is a dangerous practice, of course, especially if multiple ambiguities give rise to many combinations of possibilities. When the original source code was available, we could resolve ambiguities directly.

The second category of replication problem was gaps in the published descriptions. In two cases, published data were not complete enough to provide a rigorous test of whether distributional equivalence was achieved. In one of these cases, the author was able to provide additional data. The other gap in a published description occurred when a variable in the program could take on values of +1, 0, or -1, but was described in a way that made it appear to have only two possible values.

The third category of replication problem was situations in which the published description was clear but wrong. One example was a case in which the criteria for terminating a run of the model was not the same in the text as it was in the runs of the model for which data were reported. In another case, the description in the main text of an article was inconsistent with the appendix of the same article. Finally, there was a case in which the description in the text was a clear but inaccurate description of the model embodied in the source code.

The fourth and final category of replication problem was difficulty with the source code itself. In one case, the only source code available was from a printout so old that some of the characters were smudged beyond recognition. The last case was probably the most interesting and subtle of all. After a good deal of effort, we tracked down a difference between the original program and our reimplementation to the difference in the way two computers represented numbers. While both computers represented floating point numbers with considerable precision, they could differ in whether two numbers were exactly the same. For example, is 9/3 exactly equal to 2 + 1? In one implementation of the model it was, but in another implementation it was not. In models with nonlinear effects and path-dependence, a small difference can have a cascade of substantive effects.

BUILDING COMMUNITY

This article has discussed how to advance the art of simulation in the social sciences. It described the unique value of simulation and then offered advice for performing simulation research, focusing on the programming of a simulation model, analyzing the results, and sharing the results with others. It then discussed the importance of replicating other people's simulations and provided examples of the procedures and difficulties involved in the process of replication.

This article has already discussed suggestions for progress in methodology. The next step is to begin to establish the internal structure and boundaries of the field. In particular, converging on commonly accepted terminology would be very helpful. A host of terms is now used to describe the field. Examples are artificial society, complex system, agent-based model, multi-agent model, individual-based model, bottom-up model, and adaptive system. Having commonly accepted distinctions for these terms could certainly help specify and communicate what simulation is about.

A shared sense of the internal structure and boundaries of the field is also needed. For example, simulation in the social sciences might continue to develop primarily within the separate disciplines of economics, political science, and sociology. There are powerful forces supporting disciplinary research, including the established patterns of professional education, hiring, publication, and promotion. Nevertheless, if simulation is to realize its full potential, there must be substantial interaction across the traditional disciplines.

Progress requires the development of an interdisciplinary community of social scientists who use simulation. Progress also requires the development of an even broader community of researchers from all fields who are interested in the simulation of any kind of system with many agents. Certainly, ecology and evolutionary biology have a great deal to offer the study of decentralized adaptive systems. Likewise, computer science has recently started to pay a great deal of attention to how large systems of more or less independent artificial agents can work with each other in vast networks. And mathematics has developed some very powerful tools for the analysis of dynamic systems. Even the field of artificial life offers many insights into the vast potential of complex adaptive systems. Conversely, social scientists have a great deal to offer evolutionary biologists, computer scientists, and others because of their experience in the analysis of social systems with large numbers of interacting agents.

There are a variety of institutional arrangements that will facilitate the development of these two communities of simulators. These arrangements include journals devoted to simulation, professional organizations, conference series, funding programs, university courses, review articles, central Web sites, e-mail discussion groups, textbooks, and shared standards of research practice. Early examples of these institutional arrangements already exist. Realizing the full potential of computer simulation will require the development of these institutional arrangements for community building. Who should be better able to build new institutions than the researchers who use simulation to study real and potential societies?

ACKNOWLEDGMENTS

This article is adapted from a chapter of the same title that originally appeared in R. Conte, R. Hegselmann, and P. Terna (Eds.), *Simulating Social Phenomena,* Heidelberg, Germany, Springer-Verlag, 1997, and is reprinted with permission. I am pleased to acknowledge the help of Ted Belding, Michael Cohen, and Rick Riolo. For financial assistance, I thank Intel Corporation, the Advanced Project Research Agency through a grant to the Santa Fe Institute, and the University of Michigan LS&A College Enrichment Fund. Several paragraphs of this article have been adapted from [10] and are reprinted with permission of Princeton University Press.

REFERENCES

1. P. Bratley, B. Fox, and L. Schrage: A guide to simulation (2nd ed.). Springer-Verlag, New York, 1987.
2. W. Poundstone: The recursive universe. Contemporary Books, Chicago, 1985.
3. T. Schelling: On the ecology of micromotives. The corporate society, Robert Morris (Ed.) 1974.
4. T. Schelling: Micromotives and macrobehavior. W. W. Norton, New York, 1978.
5. H. A. Simon: A behavioral model of rational choice. Quarterly journal of economics. 69: 99-118, 1955.
6. J. G. March: Bounded rationality, ambiguity and the engineering of choice. Bell journal of economics. 9: 587-608, 1978.
7. R. Axelrod: The dissemination of culture: a model with local convergence and global polarization. J conflict resolution. 41: 203-26, 1997.
8. J. Epstein and R. Axtell: Growing artificial societies: social science from the bottom up. Washington, D.C.: MIT Press, Brookings and Cambridge, MA, 1996.
9. R. Axtell, R. Axelrod, J. Epstein, and M. D. Cohen: Aligning simulation models: a case study and results. Computational and mathematical organization theory. 1: 123-141, 1996.
10. R. Axelrod: The complexity of cooperation: agent-based models of competition and collaboration. Princeton University Press, Princeton, NJ. 1997.