Kenneth L. Judd and Leigh Tesfatsion (Eds.), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics,* Handbooks in Economics Series, North-Holland, forthcoming.

March 14, 2005

AGENT-BASED MODELING AS A BRIDGE BETWEEN DISCIPLINES

ROBERT AXELROD*
Gerald R. Ford School of Public Policy, University of Michigan

I dedicate this chapter to William Hamilton, an outstanding collaborator and a wonderful human being. Bill Hamilton literally gave his life for science. In 2000, despite the risks, he went to the jungles of central Africa to gather evidence needed to test a theory about the origin of AIDS. He contracted a virulent form of malaria that proved fatal.

Contents

---

* Acknowledgements: For their help I thank Ross Hammond, Kenneth L. Judd, Leigh Tesfatsion, Robert Marks and Tarah M. Wheeler. For financial support I thank the National Science Foundation, the Intel Corporation, and the University of Michigan LS&A College Enrichment Fund. For permission to quote William Hamilton (2002) I thank Oxford University Press.

Abstract

Using the author's own experiences, this chapter shows how agent-based modeling (ABM) can address research questions common to many disciplines, facilitate interdisciplinary collaboration, provide a useful multidisciplinary tool when the math is intractable, and reveal unity across disciplines. While ABM can be a hard sell, convergence within the agent-based community can enhance the interdisciplinary value of the methodology.

Keywords

interdisciplinary research, agent-based models, evolutionary biology, Prisoner's Dilemma

JEL classification: A12, C63, C73

1. Introduction

This chapter describes some of my experiences with agent-based modeling (ABM) as a bridge between disciplines. I offer these experiences to provide concrete examples of how agent-based modeling can help overcome the somewhat arbitrary boundaries between disciplines. I do not claim that my experiences are typical, or that my style would work well for others.

Although I am occasionally mistaken for an economist, my PhD is in political science. In graduate school, I took the micro-macro sequence designed to socialize the economic doctoral students into their discipline. I distinctly remember an occasion when the professor - a future Nobel Prize winner - was presenting a formal model of consumer behavior. A student remarked, “But that’s not how people behave.” The professor replied simply, “You’re right,” and without another word, turned back to the blackboard and continued his presentation of the model. We all got the idea.

I have undertaken ABM projects that both draw on and contribute to economics. Although I often work alone, or with a graduate student, I have also collaborated on ABM projects with political scientists, evolutionary biologists, computer scientists, and economists. This chapter draws on those experiences.

This chapter is organized in terms of five propositions, followed by some suggestions. First, here are the propositions and the research projects that I will use to illustrate them.

<u>1. ABM can address certain problems that are fundamental to many disciplines.</u> To illustrate how agent-based modeling can address a fundamental problem, I will use my computer tournaments for the iterated Prisoner’s Dilemma that addressed the question of what it takes for egoists to cooperate with each other. I will describe how computer tournaments originated from a link between game theory and artificial intelligence, how the entries drew on the strategic understanding of theorists from many disciplines, and how the results had a wide range of applications in the social sciences and beyond.

<u>2. ABM facilitates interdisciplinary collaboration.</u> An informative example is the interdisciplinary collaboration between a political scientist and an evolutionary biologist, namely myself and William Hamilton. Using Hamilton’s published memoirs, I will be able to compare and contrast his perspective with mine on how our relationship got started, and what happened along the way.

<u>3. ABM provides a useful multidisciplinary tool when the math is intractable.</u> My second collaboration with William Hamilton began when he told me about his explanation for one of the most important evolutionary puzzles: why do almost all large animals and plants reproduce sexually even at the cost of allowing only half the adults to have offspring? Hamilton had a highly original explanation, but had been unable to demonstrate its plausibility. I showed Bill how agent-based modeling could easily simulate the evolutionary effects of a dozen or more genes. Together we were able to

build an agent-based model that demonstrated that Bill's theory was indeed a biologically plausible explanation for the origin and maintenance of sexual reproduction.

4. ABM can reveal unity across disciplines. My example of this point is an agent-based model I designed based on the principle that when possible, agents will tend to align themselves into groups that are self-organized and minimize the "stress" each agent faces in its relationships with the each of the others. Scott Bennett, then a graduate student in political science, and I developed, operationalized and validated this model with alignments among the seventeen European nations that participated in World War II. Two economists, Will Mitchell and Robert E. Thomas, at my university's Business School and their graduate student, Erhard Bruderer, heard me present this theory, and immediately suggested we work together to apply it to a specific example of computing business alliances. We found that the agent-based model about military alignments could also successfully predict strategic alignments of computer companies.

5. ABM can be a hard sell. Since most formal theorists equate models with mathematical models, it is not surprising that some of them are hard to convince about the appropriateness and value of an agent-based simulation. This point is demonstrated by the kind of objections that Bill Hamilton and I met when we tried to publish what we thought were compelling results from our simulation of his evolutionary theory.

This chapter concludes with some suggestions for enhancing the interdisciplinary value of agent-based modeling.

From my perspective, agent-based modeling is not only a valuable technique for exploring models that are not mathematically tractable; it is also a wonderful way to study problems that bridge disciplinary boundaries.

2. ABM can address fundamental problems seen in many disciplines.

Agent-based modeling has proven helpful in exploring issues that arise in two or more disciplines. Examples of such issues are path dependency, the effects of adaptive versus rational behavior, the consequences of heterogeneity among agents, the design of institutional mechanisms to achieve specific goals in a population of autonomous agents, the effects of network structure, cooperation among egoists, provision of collective goods, the diffusion of innovation, and the tradeoff between exploiting current best practice and exploring for new knowledge.

My own experience includes work on the possibility of cooperation among egoists. My work on computer tournaments for the iterated Prisoner's Dilemma, for example, drew upon strategic ideas from the different disciplines of the entrants, including economics, political science, psychology, sociology, and mathematics. Simulation results and my related mathematical theorems then proved applicable to an even wider range of disciplines, as illustrated in some of my own subsequent work and that of many others.

But where did the idea for a computer tournament come from? In retrospect, I realize that it came from my interest in artificial intelligence that started while I was in high school, and an interest in game theory that started in college. In high school, a just-published article I came across fascinated me with its description of a checker playing program that learned to improve its own play (Samuel 1959). Afterwards, I followed the development of computer chess through the 1960s, and as well as the computer chess tournaments that began in 1970.

In college, I was a math major with a growing interest in international politics and especially the risk of nuclear war. While studying a then-standard text (Luce and Raiffa, 1957), I came across the iterated Prisoner's Dilemma. To me, the Prisoner's Dilemma seemed to capture the essence of the tension between doing what is good for the individual (a selfish defection) and what is good for everyone (a cooperative choice). In graduate school, while pursuing a PhD in political science, I read the intriguing research on how human subjects played the game, and how game theorists were still arguing with each other about the best way to play the game.

The literature on the iterated Prisoner's Dilemma left me somewhat frustrated because there was no clear answer to the question of how to avoid conflict, or even how an individual (or country) should play the game for itself.[1] Apparently, my frustration stayed with me because I started thinking about the problem again a dozen years later. This time I came up with the tournament idea as a means of studying these questions.

I somehow put two and two together, and realized that a good way to find a successful strategy for the iterated Prisoner's Dilemma was to hold a tournament and see what strategy would win. While I could not have articulated it then, my interest in finding out how sophisticated individuals would play to maximize their own score was probably based on the implicit belief that one would then be able to learn about the conditions under which even egoistic players would choose to cooperate.

I solicited entries from both game theorists and amateurs. Using computer chess programs as a guide, I expected the most successful strategy would have to take into account a wide variety of considerations and hence be very complicated. I was surprised

---

[1] I learned about the "solution" by backwards induction that says it pays both players to defect on the last move, hence on the next to last move, and so forth right to the beginning of the game. To me, the foresight required for a long backwards induction does not seem very realistic. I doubted that even if someone understood this logic, he or she would expect the other player to understand it, and hence the fully rational reasoning might not apply to real people. In fact, in my first computer tournament, it was common knowledge that the game would be exactly 200 moves, and I provided the entrants with an excerpt from Luce and Raiffa (1957) on backward induction in the finite iterated Prisoner's Dilemma. Although I did not realize it at the time, this design provided a test of what sophisticated and well informed researchers would expect of each other. In the event, the strategy that always defected was not entered, and three of the entrants submitted strategies that automatically defected on the last three moves, apparently because they predicted that others would do backwards induction for only *two* moves. In any case, reasoning by backwards indication does not apply when the players do not know when the iterated game will end, as was true for the second round of the tournament.

when Tit for Tat, the simplest of all the submitted strategies, and one of the simplest *possible* strategies, won the tournament.[2]

Since entries came from professors of economics, political science, psychology, sociology, and mathematics, the tournament itself illustrates how ABM can provide the means to bridge disciplines. These were mostly people who had published treatments of the iterated Prisoner's Dilemma in their own disciplinary journals. The tournament provided a way for their strategic ideas to be evaluated in the common setting, namely the rich environment that they would provide for each other. Among the most interesting results was that Tit for Tat, a strategy that could never score better than the other guy it was playing with, nonetheless won both tournaments. Wanting to reach people in many fields, I published the results in an interdisciplinary journal, the *Journal of Conflict Resolution* (Axelrod, 1980a and 1980b).

Seeing that Tit for Tat was quite robust, I used my math background to formulate and test a series of theorems about the conditions under which cooperation based upon reciprocity can emerge in a population of egoists, and then resist invasion by mutant strategies. This time I aimed for my major reference group by publishing the theorems in a disciplinary journal, *The American Political Science Review* (Axelrod, 1981).

What happened next was quite fortuitous. Following my usual practice of scanning a wide range of journals, I saw a review in a sociology journal of a fascinating study based on soldiers' diaries from World War I. The book focused on the "live and let live" system that spontaneously arose between the two sides fighting each other in the trenches. Upon reading the book, I realized that this example, in all its richness, was an apt illustration of my theory about when and how cooperation among egoists can emerge. What made the example so useful was that it showed cooperation where you might least expect it, between opposite sides in the midst of a brutal war. Yet, when viewed from the perspective I was proposing, the cooperation in trench warfare made perfect sense. When I came across this wonderful case, I thought I just might be able to write a book that could speak to a wide audience.[3]

I had no trouble finding illustrations from a wide range of fields. For instance, in economics, issues of cooperation among egoists arose in battles over barriers to trade, attainment of microcredit for those without tangible assets, strategic alliances between business, and the possibilities for tacit cooperation in a duopoly.

Seeing that the results of my computer tournaments, and the related theorems that I had provided, could address a very wide range of problems, I started to think about writing a book. When I discovered that people with little or no social science training were able to

---

[2] See Axelrod (1984). Now that I have better understanding of the effects of errors in perception and implementation, I would recommend adding a little generosity or contrition to a strategy of strict reciprocity. See Wu and Axelrod (1995).

[3] The readers of *The Evolution of Cooperation* often found the trench warfare case to be the most persuasive part of the book. I am still pondering why a single case study can be more persuasive than quantitative analysis, proofs of theorems, or a host of diverse illustrations.

understand the basic theory and the trench warfare example, I decided to try it. It took me over a year to transform and extend four of my journal articles and some new research into a book that had a chance of being read with interest by scholars and graduate students in different disciplines, and perhaps even by undergraduates and some members of the educated public. Based on its sales and citations, *The Evolution of Cooperation* succeeded beyond my hopes. From my perspective, agent-based modeling, as exemplified by the computer tournaments for the iterated Prisoner's Dilemma, demonstrated its ability to illuminate fundamental questions of interest far beyond any single discipline such as economics or political science.

One conclusion I drew from this and similar experiences was that following my own interests regardless of where they led could occasionally be not only fun, but also productive. I also realized that three of the fields that have been especially helpful to me - evolutionary biology, artificial intelligence, and game theory - I had studied on my own. I now suggest to graduate students that they should never let coursework interfere with their education.

3. ABM facilitates interdisciplinary collaboration.

I did not feel the need for a collaborator to conduct and analyze the tournaments. Nor did I feel the need for a collaborator to prove some general theorems about how cooperation based on reciprocity could get started and could resist invasion.[4] However, when I thought about the potential implications for evolutionary biology, I knew I was in over my head. I wrote to an entrant in one of my tournaments who happened to be a well-known evolutionary biologist, Oxford professor Richard Dawkins. He pointed me to another evolutionary biologist who happened to be at my own university, William Hamilton. I already knew of Hamilton's very influential theory of inclusive fitness.[5] So I gave him a call.

In his memoirs, Bill describes his reactions to this phone call.[6]

> One day in the Museum of Zoology at Ann Arbor there came a phone call from a stranger asking what I knew about evolutionarily stable strategies and for some guidance to relevant literature. (p. 118)
>
> Now on the phone to me was someone out of political science who seemed to have just the sort of idea I needed. A live games theorist was here on my own campus! Nervously, and rather the way a naturalist might hope to see his first

---

[4] See Axelrod (1981).

[5] Dawkins himself had written a lucid exposition of Hamilton's theory of inclusive fitness in a book entitled *The Selfish Gene* (1976, 1989). Once you read this book, you will see why your genes can be considered "selfish" and how your selfish genes use you to get themselves reproduced – but not necessarily in ways that are to *your* advantage. Spooky.

[6] All quotes are from William D. Hamilton (2002).

> mountain lion in the woods, I had long yearned for and dreaded an encounter with a games theorist. How did they think? What were their dens full of? Axelrod on the phone sounded nice and, very surprising to me, he was more than a bit biological in his manner of thinking. I sensed at once a possibility that the real games theorists might be going to turn out to be a kind of kindred to us [biologists].(p. 120)

Had Bill known of my long-standing interest in evolutionary theory, he might not have been quite so surprised that my thinking was more than a bit biological. For example, in high school I wrote a computer simulation to study hypothetical life forms and environments.[7] This early interest in evolution was nurtured during college by a summer at the University of Chicago's Committee on Mathematical Biology.

That first phone call led to a lunch where he proposed that we work together.

> Soon after the lunch again I proposed that the work seemed so interesting biologically we might try writing it up for a joint paper in *Science;* [Axelrod's] contribution would be the basic ideas plus the description of his tournaments, and mine to add a natural scientist's style and some biological illustrations. (p. 122)

I was delighted to accept Bill's invitation to collaborate. Despite coming from different disciplines, Bill and I shared not only mathematical training, but also a desire to get at the heart of things. Bill had even published some work on the Prisoner's Dilemma, although he was hoping to get away from that when I dragged him back.

Bill's proposed division of labor turned out to be a reasonable description of how the collaboration developed. I gradually realized, however, just how much was included by Bill's modest formulation of adding "a natural scientist's style and some biological illustrations." Bill's naturalist's style included having at his fingertips an astonishing knowledge of species from bacteria to primates. His knowledge would be equivalent to an economist knowing much of what there is to know about hundreds, if not thousands, of companies of every type from GM and Microsoft to a self-employed sidewalk vendor.

His experience as a naturalist often gave him the capacity to check out the plausibility of an idea with pertinent examples right off the top of his head. It also helped him to generate surprising new ideas.

---

[7] The recognition I received for this work from the Westinghouse (now Intel) Science Talent Search contributed to my readiness ever since to follow my own instincts and not worry about what was in the mainstream of any particular discipline. The simulation had agents who responded to their environment, but was not an agent-based model because the agents neither interacted with each other, nor changed over time. In 1960-61, Northwestern University gave me some time on their one and only machine, an IBM 650 the size of four refrigerators. It had only 20k memory – about a millionth of my current laptop's memory.

Bill's disciplinary training as an evolutionary biologist and a naturalist proved essential to making our theoretical work compelling to biologists. He was to identify and exploit pertinent biological examples so that biologists could see what we were talking about. While not all of his proposed applications have been borne out, he was able to demonstrate the potential relevance of agent-based computer tournaments for the major biological puzzle of why individuals cooperate with unrelated others.[8] Second, he was able to explain what our contribution added to what was already understood about evolution. Specifically, he showed how our modeling work provides a solid foundation for many of the insights about altruism formulated years earlier by Robert Trivers (1971). Bill was also able to show how our model could be used by other evolutionary biologists to formulate and test new hypotheses about animal behavior,[9] as well as explore dozens of variants of the simple iterated Prisoner's Dilemma.[10]

At the beginning, Bill and I took a while to get used to each other's style. For example, when I asked Bill a question he sometimes thought long and hard before saying a word. His face took on a blank look, his gaze was in the distance, and I could almost hear the wheels spinning inside his head for the longest time. I learned to be patient. When he finally spoke, it would be either a true insight, or casual remark on a totally different subject.

Here is how Bill saw us working together.

> That brilliant cartoonist of the journal *American Scientist*, Sidney Harris, has a picture where a mathematician covers the blackboard with an outpouring of his formal demonstration. … [I]t starts top left on the blackboard and ends bottom right with a triumphant 'QED'. Halfway down, though, one sees a gap in the stream where is written in plain English: 'Then a miracle occurs', after which the mathematical argument goes on. Chalk still in his hand, the author of this *quod est demonstrandum* now stands back and watches with a cold dislike an elderly mathematician who peers at the words in the gap and says: 'But I think you need to be a bit more explicit-here in step two.' I easily imagine myself to be that enthusiast with the chalk and I also think of many castings for the elderly critic. Yet how easy it is to imagine a third figure-Bob-in the background of this picture, saying cheerfully: 'But maybe he has something all the same, maybe that piece can be fixed up. What if. ..' (p. 123)

I shared Bill's surprise at how well we worked together. As he put it,

---

[8] Bill was already well known for his rigorous treatment of how evolution might cause an individual to be altruistic toward a close relative (Hamilton, 1964). Because Hamilton showed how to treat the unit of selection as the gene rather than the individual to the gene, this work has been called "the only true advance since Darwin in our understanding of natural selection" (Trivers, 2000).

[9] For early confirmations in bats, fish and primates, as well as early extensions of the iterated Prisoner's Dilemma framework, see the sources cited in Axelrod and Dion (1988). Recently, even viruses have been found to play the Prisoner's Dilemma.(Turner and Chao, 1999)**.**

[10] For early develops of the theoretical framework see Axelrod and Dion (1988). For a twenty year retrospective, see Hoffman (2000).

> I would have thought it a leg-pull at the time if someone had told me of a future when I would find it more rewarding to talk 'patterns' to political scientists rather than to fellow biologists. (p. 126)

Perhaps the most important thing we shared was our aesthetic sense.

> [A]n intuitive understanding between us was immediate. Both of us always liked to be always understanding new things and to be listening more than talking; both of us had little inclination for the social manoeuvring, all the 'who-should-bow-lowest' stuff, which so often wastes time and adrenalin as new social intercourse starts. Bob is the more logical, but beyond this what we certainly share strongly is a sense for a hard-to-define aesthetic grace that may lurk in a proposition, that which makes one want to believe it before any proof and in the midst a confusion and even antagonism of details. Such grace in an idea seems often to mean that it is right. Rather as I have a quasi-professional artist as my maternal grandmother, Bob has one closer to him-his father. Such forebears perhaps give to both of us the streak that judges claims not in isolation but rather by the shapes that may come to be formed from their interlock, rather as brush strokes in a painting, shapeless or even misplaced considered individually, are overlooked as they join to create a whole…. (p. 122)

I see a further connection between art and modeling. My father painted to express how he saw the world that day, highlighting what was important to him by leaving out what was not. Likewise, I see my modeling, especially my agent-based modeling, as an expression of how I see some social dynamic, highlighting what I regard as important, and leaving out everything else.

Our differing disciplinary backgrounds would show up in surprising ways, such our reactions to visiting the church where Shakespeare is buried. I pondered the social science question of why Shakespeare might have wanted others to read a mediocre poem on his gravestone, and Bill pondered the biological puzzle of why a very rare plant was growing on the fence outside the church.

> Anyway, between us and with surprisingly little difficulty we pushed our paper into *Science.*[11] Once published it drew so much interest that it won us the Newcomb-Cleveland Prize as *Science's* supposed best paper for its year, 1981. (p. 123-4)

4. ABM provides a useful multidisciplinary tool when the math is intractable.

---

[11] Axelrod and Hamilton (1981).

Agent-based modeling can also be useful for discovering regularities that might suggest theorems that can then be proved. For example, my finding that Tit for Tat did well playing with a wide variety of other strategies, led me to expect that something very general could be proved about the conditions under which Tit for Tat could withstand "invasion" by *any* other strategy. And so it turned out.[12]

A second collaboration with Bill Hamilton demonstrates another valuable characteristic of agent-based modeling: its ability to analyze problems by simulation even when mathematical analysis is impossible.

When Bill took a very prestigious position at Oxford, we still kept in touch. We shared our on-going thinking. One day, about five or six years after our first collaboration was finished, Bill told me about a truly amazing theory he was developing. The theory proposed an answer to one of biology's largest unresolved puzzles: why have most large animals and plants evolved to reproduce sexually? The reason this is such a puzzle is that sexual reproduction has a huge cost: only half the population has offspring. What might be the advantage of sexual reproduction that is so great that it can overcome this two-fold cost compared to asexual reproduction?

There was already a serious contender whose leading advocate was the Russian geneticist Alexei Kondrashov. Kondrashov's explanation was based on the possibility that sexual reproduction might be helpful for bearing the cumulative burden of many generations of deleterious mutations. Bill's theory was completely different. Put simply, he thought of sexual reproduction as an adaptation to resist parasites.[13] This struck me as a totally bizarre, but intriguing idea.[14]

Bill explained to me that there was a serious problem with convincing others that his theory could, in fact, account for the two-for-one burden of sexual reproduction. The problem was that the equations that described the process were totally intractable when the genetic markers had more than more than two or three loci. Yet, the whole idea relies on there being many loci so that it would not be trivial for the parasites to match them. When I heard this, I responded to Bill with something like, "No problem. I know a

---

[12] See Proposition 2 in Axelrod (1984, pp. 207-9). Taylor (1976) had already proved that Tit for Tat could resist invasion by several specific strategies, but the success of Tit for Tat with a wide range other strategies in the tournaments suggested to me that it was worthwhile to seek a theorem that would apply to *all* other strategies. By viewing Tit for Tat as a two-state finite automaton, I was able to prove such a general result.

[13] Bill liked this formulation of mine, and we used it as the title of our paper.

[14] Bill's reasoning was that parasites are ubiquitous, and their short life spans give them the advantage of being able to adapt quickly to an ever-changing host population. If the host population reproduced *asexually*, a line of parasites that had evolved to mimic the genetic markers on the cells of one host would automatically be well adapted to mimic the genetic markers of its offspring. On the other hand, if the hosts reproduced sexually, their offspring would not be virtual carbon copies of either of their parents, and thus would not be as vulnerable to a line of parasites that had become adapted to match the genetic loci of one parent or the other.

method to simulate the evolution of populations with a lot of genetic markers. The method is called the Genetic Algorithm, and I've already used it to simulate a population of individuals each of whom has seventy genes."[15]

I explained to Bill that a computer scientist, John Holland, had been inspired by the success of biological evolution in finding "solutions" to difficult problems by means of competition among an evolving population of agents.[16] Based on the evolutionary analogue, including the possibility for sexual reproduction, Holland developed the Genetic Algorithm as an artificial intelligence technique. I could simply turn this technique around and help Bill simulate biological evolution, with or without sex. Since Bill was used to thinking in terms of heterogeneous populations of autonomous individuals, he readily grasped the idea of agent-based modeling. He also grasped without difficulty that an agent-based simulation was capable of demonstrating that certain assumptions are sufficient to generate certain results, even if the same results could not be proven mathematically.[17]

So, working with a computer science graduate student, Reiko Tanese, we built an agent-based model with two co-evolving populations: hosts with long life spans, and parasites with short life spans. If a parasite interacted with a host of similar marker genes, it killed the host and reproduced. In the simulation, the parasite population would tend to evolve to concentrate in the region of the "genetic space" where there were many hosts. Thus, successful hosts tended to suffer from increasing numbers of deadly parasites, reducing the numbers of those hosts. Meanwhile, other types of hosts with very different genetic markers might thrive. Then the process would repeat itself as the population of parasites tracked the ever-changing population of hosts. The system would always be out of equilibrium.[18]

Bill was pleased with the results of our agent-based simulations. He felt that

> the notion I had started with, that even against sex's full halving inefficiency the problem could be solved by looking at the need of a population to manoeuvre against its many rapidly evolving parasites, with these differentiating resistance tendencies at many host loci (the more the better), had been vindicated.(p. 561)

---

[15] I used this evolutionary technique to avoid having to run new tournaments indefinitely. See Axelrod (1987). I had earlier developed a technique now known as replicator dynamics, to study an interacting population with many different types of individuals, but without any mutation to introduce new types (Axelrod, 1980b, and Axelrod, 1984, pp. 48-54).

[16] See Holland (1975 and 1992), and Riolo (1992).

[17] Not proven by humans at least. Epstein, in this volume, points out that the premises of simulations can themselves be regarded as mathematical statements, and results as deductions derived from those statements.

[18] Agent-based models are convenient for studying out-of-equilibrium dynamics. Real economies may be perpetually out of equilibrium, for instance if there is continual innovation (Nelson and Winter 1982). Systems far from equilibrium are notoriously difficult to analyze mathematically, and perhaps for that reason are often downplayed in neoclassical economics. Agent-based modeling allows the analysis systems that are far-from-equilibrium

> Returning to the story of the work, once Reiko under Bob's guidance had done the program, I experimented with it by e-mailing her or Bob with requests for chosen runs. At one point I visited the University of Michigan at Ann Arbor and worked for a fortnight intensively on modifications to the program with Reiko- this came after a bad patch of misunderstandings and unpromising runs that had caused us all to become somewhat pessimistic. (p. 606)

It seems to me that agent-based modeling is quite vulnerable to misunderstanding, even among the collaborators themselves. In our case the problem arose while we were exploring different ways to model host-parasite interaction. At one point Bill sent an e-mail from Oxford asking Reiko and me to undo our recent changes and try something else that he described. It wasn't until a month or so later that Bill noticed the unpromising runs might be caused by our simulation program not doing quite what he had in mind. We eventually traced the problem to a misunderstanding between us about whether Bill's request to remove our recent changes referred to the previous day's work, or the previous week's work.[19]

> Daily Reiko sprinkled me and Bob, like tender house plants, with her floppy disks bearing her updated codes….(p. 607)

> Our model had achieved results that others had stated impossible with the tools we were allowing ourselves. Many of the dragons that had oppressed individual-advantage models in the past seemed to us to be slain. …[O]ur explicit modeling of a large number of loci in a Red Queen situation[20] certainly was [new] and the increase of stability of sex that came with the growth of numbers of loci made the most dramatic feature in our results. (p. 602)

> It is the paper that I regard as containing the second most important of all my contributions to evolution theory.[21] That second joint paper of 1990 (actually mainly written some three or so years before) was to be the first model where sex proved itself able to beat any asex competitor immediately and under very widely plausible assumptions. (p. 560)

---

[19] While a mathematical proof can usually be checked for accuracy without great difficulty, the same can not be said for an agent-based simulation program. The frequency of this problem became evident when I was part of a team that tried to replicate the results from the published description of the eight agent-based models (Axtell et al. 1996). We found that in most cases it was not easy. In one case, it took us about four months to track the problem to an inconsistency between the published account and the actual code used to implement it. Results from macro-economic models are also notoriously difficult to replicate from published descriptions, even when the identical data set is used.

[20] Bill is referring here to the character in *Alice Through the Looking Glass* who says, "It takes all the running you can do, to keep in the same place."

[21] Hamilton, Axelrod and Tanese (1990). Bill regarded his most important paper to be the one that presented his formal theory of inclusive fitness (Hamilton, 1964). As mentioned earlier, Dawkins (1976; 1989) provides a lucid exposition of Hamilton's theory of inclusive fitness.

5. ABM can reveal unity across disciplines.

So far, I have described my experiences of using agent-based modeling to bridge disciplines by addressing fundamental problems, by facilitating collaboration, and by avoiding intractable mathematics. Finally, I want to discuss an example of how an agent-based model designed for a specific problem in one discipline can sometimes be applied directly to an apparently quite different problem in another discipline.

My own specialty in political science is international security affairs. I wanted to predict alignments in war. I did not want to beg the question by taking into account any alliances the countries might have already formed. The problem is exemplified by the mutual hatred on the eve of World War II between Germany, Britain, and the Soviet Union. If they all hated each other, what would predict their alignment into just two opposing sides when the war came?

The model began by assigning countries at random to one of two sides, and giving each country, one at a time, the opportunity to change sides if it would reduce the "stress" of their being aligned with countries they were repelled from and/or *not* aligned with countries they were attracted to. Naturally, the felt stress would also have taken into account the relative importance of each of the other countries. Scott Bennett, a political science graduate student, and I operationalized the pairwise propensities by combining five previously identified factors causing attraction or repulsion. These factors included things like shared religion and border disputes. We operationalized the importance of each country by the magnitude of its relative strength at the time. We then simulated the process using the seventeen European countries that became involved in World War II. No matter which of the 65,536 different alignments we started with, the agents always organized themselves into one of just two alignments. One of these two is almost exactly what happened in World War II.[22]

Around this time, I was invited to present my latest research at Michigan's Business School. After my talk, two economists from the school, Will Mitchell and Robert E. Thomas, came up to Scott and myself. They told us that our work reminded them of the business coalitions that often form to compete over whose preferred standard will dominate an industry. They said they had in mind the specific case in which eight computer companies joined one of the two coalitions that competed over which version of the UNIX operating system would prevail. We decided to see if we could account for the specific alignment of companies in the UNIX case. We used exactly the same theory, and simply adapted the measures of pairwise propensity and relative size so they made sense for the UNIX case. For example, we assumed that a company would find it more stressful to align with a company that was largely in the same market as it was, compared

---

[22] See Axelrod and Bennett (1993). The one mistaken prediction was that Portugal with its fascist government would side with Germany and Italy, but it actually stayed with its long-term ally, Britain. The other alignment was essentially a pro-vs.-anti Communist alignment. On another point, economists sometimes ask me why the agents in this model might not keep switching sides forever. The short answer is that the pairwise propensities are symmetric, so "stress" provides a Lyapunov function.

to a company that was mainly in a completely different market. We were delighted to find that the agent-based model of military alignments was also successful at predicting the pattern of strategic alignments among the eight computer companies involved in the UNIX case.[23]

6. ABM can be a hard sell.

As noted earlier, my collaboration with Bill Hamilton on cooperation in biological systems was accepted for publication with little problem. Just the opposite was true of our second collaboration. Our simulation of Bill's theory that sexual reproduction could be an adaptation to resist parasites had a hard time getting published.

First, we tried *Nature*, a leading scientific journal closely followed by biologists of all types. The referees had many complaints, chiefly about the robustness of our results. So, we did many more runs under a broad range of parameters to show that the explanation worked under a wide range of realistic conditions. We thought our second try had nailed our point.

> Nevertheless when the revised paper went back to the referees with these new experiments included, but with no change to our centralizing of the Homo-like life history, we found all our new points left uncommented and the manuscript rejected by the referees even more curtly than before. Two of them indeed dug out new objections they hadn't thought of [the] first time and claimed to see no substantial changes in the rest….(p. 608)

After our revised version was rejected at *Nature*, we submitted our paper to *Science*, another leading scientific journal widely read by biologists. We were also rejected by *Science*, which left us a little dejected.

> Failing with these I sent it in preliminary way to an editor of the Royal Society journals to see if they would be interested, but the comments I received were as discouraging as the rest. It particularly shamed me to have to tell Bob that even the society that supported me in general believed me to be over the hill on this topic… (p. 609)
>
> One of the puzzles about the dislike, even contempt, the work … seemed to arouse in my evolutionary peers is that it was as if we had been unable to explain what we were thinking. …And yet while one referee praised our style, another described the paper as written very badly; because neither said anything good about the ideas or content I presume that even the one that liked the writing found it a kind of eloquent twittering. (p. 601-2)

---

[23] See Axelrod et al. (1995).

> [T]he only intelligible claim in [one review] was that we had not reported on any simulations outside the range we had studied in detail. … If one criticized every paper studying some feature of one-locus population genetics, for example, on grounds that it hadn't yet probed into even just possible two-locus complications (or hadn't reported having done so), a substantial fraction of the literature of population genetics would have stayed unpublished. (p. 613).
>
> Our statement that we had tested the model much more widely than we covered in the states we reported evidently wasn't believed, as also was the case with our description of the model. Several referees said this wasn't adequate; and yet it was quite as thoroughly described as models usually are in papers whose results rely on simulation. ….[In fact] a subsequent team (Richard Ladle and Rufus Johnstone, later joined by Olivia Judson) reproduced and extended our model purely [from] the paper's specification. Ladle and Johnstone did not even tell me they were working on this until our major results had been verified. (p. 610)

Bill was surprised by the difficulties we were having.

> The above record of rejections probably actually isn't long compared with some that much more revolutionary yet valid papers have received from journals. What, for example, about the attempts of Alfred Wegener to publish on continental drift, or Ignaz Semmelweiss to publish on puerperal fever, or Richard Altmann on the symbiotic origin of the mitochondrion? On the other hand, at the time we were submitting neither Bob nor I was an unknown scientist and neither of us had a reputation for mistaken or trivial ideas. The number of suspicious and hostile referees we found had come, therefore, as a considerable surprise. (p. 609)
>
> …my efforts to remould [our simulation] to appease the latest whims of referees … never worked: the referees always had new objections; dislike for our solution seemed to be unbounded. (p. 562).

Nevertheless, we were dogged. We kept revising the paper to take into account, as best we could, the reviewers' criticisms. Finally, the fifth version was successful at the *Proceedings of the National Academy of Sciences, USA.* At last, two reviewers saw the point of our paper, and one was even enthusiastic.

Why was our agent-based model (Hamilton, Axelrod and Tanese, 1990) such a hard sell? It was not because our model was less realistic than analytic models of evolution that had already been published, or our work did not break new ground, or that the problem was not important. So what was it? Bill thought about it this way:

> Simulation in itself admittedly isn't understanding and various previous papers, including some of my own …, had already drawn attention to the kinds of possibilities we were now testing. The simpler analytical discussions and

> models, however, including again my own, all had had severe snags and none showed any chance to be general. Besides treating many loci and many parasites at once-obviously much closer to the real situation (and the importance of our studying truly many loci, not just three or four, cannot be overstated)-we had brought in a variable life history that I consider to be much more realistic than is typical in most evolutionary modelling. …(p. 603)
>
> Nor could anyone pretend that this theme of evolution of sex was a narrow one nor of specialist interest only: from Erasmus Darwin to the present time, sex has repeatedly been saluted as one of biology's supreme problems, perhaps its very greatest. Hence Bob Axelrod and I at first believed that our model, with its realism and its dramatic success under conditions others had deemed impossible for it, was virtually sure to be acceptable to one of the major general scientific journals such as *Science* or *Nature.* (p. 604-5)

We suspected that part of the problem was that the reviewers were threatened by our application of Bill's theory to the case of human-like organisms – organisms similar to the reviewers themselves. It must not have been easy for them to accept that their own sexuality derived from the selective pressure of parasites.

Since we wanted to demonstrate that his theory could explain sexual reproduction in humans, Bill thought it was important that we include the salient characteristics of human reproduction. For instance, he wanted to include the fact that humans are not fertile for the first dozen or so years of their life. I, however, wanted our model to be as simple as possible to make it easier to understand and appreciate. This was the only significant disagreement we ever had. Since it was Bill's theory and Bill's audience, I deferred to his preferences in this regard. So one reason our model might have been so hard to sell is that it included some realistic details that may have obscured the logic of the simulation. On the other hand, Bill was probably right that had we *not* included these details, the reviewers would complain that we had not demonstrated the theory could account for sexual reproduction in humans. Sometimes you just can't win.[24]

Agent-based modeling is not alone in suffering from the inevitable trade-off between realism and clarity. Analytic models of economic, political, and social phenomena must deal with the same tradeoff.

There are at least two factors, however, which make it harder to sell an agent-based model than a model that can be analyzed mathematically. The first problem is that most reviewers (and potential readers) of theoretical work are familiar with the logic of deductive mathematics, but not with logic of agent-based modeling. Indeed, they often demand that the results of an agent-based model must be as general as the results of an

---

[24] We might have tried another tactic. We could have first introduced a minimal version of the model to highlight the essential mechanisms to demonstrate that Bill's theory could, in principle, explain how sexual reproduction could overcome its two-fold cost. We could then have provided a more realistic and detailed simulation to show the theory also applied to situations characteristic of human life spans. Unfortunately, the journals we aimed for had such strict page limits that we were not able to write our paper this way.

analytic model. This point is illustrated by neoclassical economic models that rigorously demonstrate that (under certain assumptions) raising the minimum wage will lower total employment. Now suppose that an agent-based model demonstrates the same effect under less restrictive assumptions about the uniformity of the labor market, but much greater specificity about the value of the parameters describing the situation. A mathematically inclined reader is likely to want to know how robust the results are, and agent-based modeling may not be able to provide a definitive answer to that question.

A typical mathematical result might take the form “For all $A>0$ and all $B>0$, $f(A,B) > 0$” where f is some given function. An agent-based model typically needs to assume specific values for certain parameters in order for the simulation to run. The simulation might be run many times, with a range of positive values for A and B, and get the same result each time. But the reviewer can always say, “but have you tried $A=1/9$, or $B=10{,}000{,}000$? And if it works for those values, what about $A = B^2$ where B is some integer multiple of pi?” In general, there is always a question of the robustness of the results of a simulation, unless the simulation results suggest a theorem that can be proved analytically.

Even if the reviewer is satisfied with the range of parameter values that have been tested, he or she might think up some new variations of the model to inquire about. Demands to check new variants of the model as well as new parametric values in the original model can make the review process seem almost endless. What is worse is that a reviewer with a not-so-legitimate problem with the submission can always use “insufficient” checks for robustness as a cover for a negative review.

Years later, Bill noticed that other researchers doing work related to his theory of sexual reproduction met with the same problem we did. Bill put a positive spin on it.

> [I]nsinuations of unreliability [i.e. insufficient checks for robustness] so extremely similar to those being directed at my work with Bob and Reiko…tended to reassure me that our rejection didn't necessarily mean that… I'd completely lost my marbles. (p. 614)

7. Convergence within the ABM community can enhance the interdisciplinary value of ABM.

In closing, I have three suggestions to facilitate interdisciplinary work with agent-based modeling.[25]

First, the agent-based modeling community should converge on standards for testing the robustness of an agent-based model. My own experience suggests that the lack of such standards can make agent-based modeling a hard sell. Just as the social sciences have converged on 0.05 as the minimal standard of statistical significance, the agent-based

---

[25] For my suggestions on how to actually do agent-based modeling, see Axelrod (1997a).

modeling community should converge on standards appropriate to the kinds of simulations we do.[26]

Second, the agent-based modeling community should converge on its tools. Just as there is a convergence on regression as one of the standard tools of statistics in the social sciences, there should be convergence on the basic tools for agent-based modeling. This is already underway. For instance, in models with a two-dimensional space, there is already something close to a consensus that unless there is a stated reason not to, the borders of the space should wrap around, e.g. making the top row adjacent to the bottom row. On the other hand, there is less consensus on whether the default assumption about whether a given cell should have four neighbors (the cells to the north, south, east and west) or eight neighbors (those four plus the diagonal cells).[27] Greater standardization of programming tools would also be helpful.[28]

Finally, the agent-based community should converge on a set of fundamental concepts and results. Just as the content of a microeconomics course at any given level has become largely standardized, it would be helpful if the same would become possible for courses in agent-based modeling. Textbooks are one way in which this convergence could be promoted. Before then, however, candidates for the shared set of fundamentals could take the form of topics or even specific readings that anyone interested in the field could be expected to know. An example of a strong candidate for inclusion is Schelling's well-known model of residential mobility that demonstrates how an emergent property like segregation can occur even if everyone is quite tolerant (Schelling 1978).

Agent-based modelers actually know quite a bit about the possibilities for convergence in a heterogeneous population of autonomous agents – such as themselves.[29] Fortunately, the bottom-up form of convergence is already underway, but we need to be wary of convergence taking place only within rather than across disciplines.[30] As a step in this direction, Axelrod and Tesfatsion (2005) have developed a Guide to Newcomers to Agent-Based Modeling.[31] Perhaps we are approaching the time when it becomes possible to develop a more or less authoritative statement of proposed core readings and best practices. This volume itself is a major step in that direction, thereby facilitating the potential of agent-based modeling to serve as a bridge between disciplines.

---

[26] For example, one might halve and double the base values of each parameter to see if the results hold up across this wide range. John Miller has proposed another possibility (Miller 1998). He suggests searching for the largest and smallest values of each parameter that will maintain the central result. In other words, we should report what extreme values cause for the model to "break."

[27] These are called the Von Neumann neighborhood and the Moore neighborhood, respectively.

[28] This is easier said than done. An early attempt called Swarm has had limited success. With object oriented programming and sharable languages like Java, the prospects are better at least for shared libraries of commonly used procedures. Ascape is a good example.

[29] My two cents worth are models of the emergence of norms (Axelrod, 1986) and the dissemination of culture (Axelrod, 1997b).

[30] For an agent-based model of local convergence and global polarization, see Axelrod (1997b).

[31] The Guide is an Appendix to this volume, and is also available at http://www.econ.iastate.edu/tesfatsi/abmread.htm.

References

Axelrod, R. (1980a) Effective Choice in the Prisoner's Dilemma, Journal of Conflict Resolution 24:3-25. Included in revised form as part of Chapter 2 and Appendix A of Axelrod (1984).

Axelrod, R. (1980b) More Effective Choice in the Prisoner's Dilemma, Journal of Conflict Resolution 24:379-403. Included in revised form as part of Chapter 2 and Appendix A of Axelrod (1984).

Axelrod, R. (1981) Emergence of Cooperation Among Egoists, American Political Science Review 75:306-18. Included in revised form as Chapter 3 and Appendix B of Axelrod (1984).

Axelrod, R. (1984) The Evolution of Cooperation (NY: Basic Books).

Axelrod, R. (1986) An Evolutionary Approach to Norms, American Political Science Review 80:1095-1111. Included with an introduction in Axelrod (1997c).

Axelrod, R. (1987) The Evolution of Strategies in the Iterated Prisoner's Dilemma, in: L. Davis, ed., Genetic Algorithms and Simulated Annealing (London: Pitman, and Los Altos, CA: M. Kaufman, 1987), 32-41. Included with an introduction in Axelrod (1997c).

Axelrod, R. (1997a) Advancing the Art of Simulation in the Social Sciences, in: R. Conte, R. Hegselmann and P. Terna, eds., Simulating Social Phenomena (Berlin: Springer, 21-40. Included with an introduction in Axelrod (1997c). An updated version of this paper is in The Journal of the Japan Society for Management Information, Special Issue: Agent-Based Modeling, Vol. 12, No. 3, Dec. 2003, and is available at http://www-personal.umich.edu/~axe/

Axelrod, R. (1997b) The Dissemination of Culture: A Model with Local Convergence and Global Polarization, Journal of Conflict Resolution 41:203-26.

Axelrod, R. (1997c) The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration, (Princeton, NJ: Princeton University Press). Some of the chapters in this book are available at http://www-personal.umich.edu/~axe/

Axelrod, R. and D. Dion (1988) The Further Evolution of Cooperation, Science 242:1385-1390.

Axelrod, R., W. Mitchell, R. E. Thomas, D. S. Bennett, and E. Bruderer (1995) Coalition Formation in Standard-Setting Alliances, Management Science 41:1493-1508. Reprinted with introduction in Axelrod (1997c).

Axelrod, R. and M. D. Cohen. (2000) Harnessing Complexity: Organizational Implications of a Scientific Frontier (NY: Free Press).

Axelrod, R. and S. Bennett (1993) A Landscape Theory of Aggregation, British Journal of Political Science 23:211-33. Included with an introduction in Axelrod (1997c).

Axelrod, R. and W. D. Hamilton (1981) The Evolution of Cooperation, Science 211:1390-1396. Reprinted in modified form as Chapter 5 of Axelrod (1984).

Axelrod, R. and Tesfatsion (2005) A Guide for Newcomers to Agent-Based Modeling, Appendix to this volume.

Axtell, R., R. Axelrod, J. Epstein, and M. D. Cohen (1996) Aligning Simulation Models: A Case Study and Results, Computational and Mathematical Organization Theory 1:123-141.

Dawkins, R. (1976) The Selfish Gene. (Oxford: Oxford University Press). New Edition 1989.

Hamilton, W. D. (1964) The genetical evolution of social behaviour. I and II, Journal of Theoretical Biology 7:1-52.

Hamilton, W. D. (2002) Narrow Roads of Gene Land: The Collected Papers of W. D. Hamilton, vol. 2, The Evolution of Sex. (Oxford, Oxford U Press), 117-32, 561-6, and 601-15.

Hamilton W. D., R. Axelrod and R. Tanese (1990) Sexual Reproduction as an Adaptation to Resist Parasites, Proceedings of the National Academy of Sciences (USA), 87:3566-3573.

Hoffmann, R. (2000) Twenty Years on: The Evolution of Cooperation Revisited, Journal of Artificial Societies and Social Simulation 3(2). Available at <http://www.soc.surrey.ac.uk/JASSS/3/2/forum/1.html>

Holland, J. H. (1975) Adaptation in Natural and Artificial Systems: An Introductory Analysis With Applications to Biology, Control, and Artificial Intelligence. (Ann Arbor, MI: University of Michigan Press). Reissued Cambridge, Mass.: MIT Press, 1992.

Holland, J. H. (1992) Genetic Algorithms, Scientific American, July: 44ff.

Luce, R. D. and H. Raiffa (1957) Games and Decisions; Introduction and Critical Survey (New York, Wiley).

Maynard Smith, J. (1982) Evolution and the Theory of Games (Cambridge, UK: Cambridge University Press).

Miller, J. H. (1998) Active Nonlinear Tests (ANTs) of Complex Simulations Models, Management Science, 44:820-30.

Nelson, R. and S. G. Winter (1982) An Evolutionary Theory of Economic Change (Cambridge: Harvard University Press).

Riolo, R. L. (1992) Survival of the Fittest Bits, Scientific American, July: 89ff.

Samuel, A. (1959) Some Studies in Machine Learning Using the Game of Checkers, IBM Journal of Research and Development, 3(3):210-229.

Schelling, T. (1978) Micromotives and Macro_Behavior (NY: Norton).

Taylor, M. (1976) Anarchy and Cooperation (New York: Wiley).

Trivers, R. (1971) The Evolution of Reciprocal Altruism. Quarterly Journal of Biology. 46:35-57.

Trivers, R. (2000) Obituary: William Donald Hamilton (1936-2000), Nature 404:828.

Turner, P. E. and L. Chao (1999) Prisoner’s Dilemma in an RNA Virus, Nature, 398:367-8.

Wu, J. and R. Axelrod (1995) How to Cope with Noise in the Iterated Prisoner's Dilemma, Journal of Conflict Resolution, 39:183-189. Reprinted with an introduction in Axelrod (1997c).