

How Political Science Can Enrich Other Disciplines

The 2013 Skytte Prize Lecture

Robert Axelrod*

It's a great honor being here today. And I appreciate the chance to speak at this very distinguished lecture series in honor of Johann Skytte.

I understand that when Johann Skytte donated the world's first professorship in political science, he originally named it the professorship of Eloquence and Government. Well, the government part flourished, but you'll see today that the perhaps the eloquence part did not do quite as well.

It's a wonderful coincidence that almost thirty years ago the first translation of my book on the *Evolution of Cooperation* was the Swedish edition. And just this month, an Arabic edition is being published, and I hope that my own modest efforts in recent years to promote peace in the Middle East has had something to do with the new availability of my book in Arabic.

But today, I don't plan to talk about international politics. Instead, I'd like to talk about the discipline of Political Science itself. I'll do so in terms of some things Political Science has gained from other disciplines and especially what it has to offer to other disciplines.

To illustrate my themes, I'll use my own experience with research on the evolution of cooperation. In the interests of full disclosure, some pieces of what I'll be talking about today I have already published. But I think it was Paul Samuelson who said that there is only one absolute ethical principle in academics, and that is that you can't publish the same paper more than three times. By that standard, I'm quite safe.

* Robert Axelrod, Gerald R. Ford School of Public Policy and Department of Political Science
University of Michigan, Ann Arbor. E-mail: axe@umich.edu

Borrowing From Other Disciplines

As you probably know, my work on cooperation borrows a lot from game theory, and in particular the iterated Prisoner's Dilemma game. My original interest in game theory arose from a concern about what the United Nations Charter calls the 'scourge of war.' My work focused on how to achieve cooperation among egoists, that is to say among people or countries who care about themselves, and only care about others to the extent it can help *them*.

The iterated Prisoner's Dilemma seemed to me 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).

My driving question was, 'Under what conditions can egoists cooperate with each other, when there is no central authority to police their interaction.'

Because I analyzed this question at a very general level, my findings had wide applicability not only in political science, but also economics, sociology, anthropology, evolutionary biology, and many other fields that study how people (or even biological cells) interact with each other. Even better, there began to be a conversation across all of these fields. In this sense, game theory approaches to problems of cooperation not only helped promote work across many disciplines; it also helped erase boundaries between disciplines. It did this by providing a common set of concepts, questions, mathematical theories, simulation results, and empirical findings.

When I started my work on the iterated Prisoner's Dilemma with the question of when egoists will cooperate with each other, I realized I first needed to answer another question, namely what's a good strategy to use in the iterated Prisoner's Dilemma? Surprisingly, there were a variety of proposed answers in the literature, but no definitive answer since it all appeared to depend on what strategy the other player in the game was using.

Therefore I was intrigued by how to choose among the many strategies that had been proposed to play this game effectively. To answer this question, I imported ideas from economics, artificial intelligence and evolutionary biology. Let me tell you about that.

Economics was the discipline where most of game theory was developed. My formal education in economics included a graduate course designed not for political scientists but for PhD students in economics. It was taught by a future Nobel Prize Winner. One day while the teacher was putting on the blackboard a rational choice model of consumer behavior, a student interrupted to say, 'But that's not how people behave.' The instructor responded by saying, 'You are right,' and without saying another word, turned back to the blackboard and continued his exposition. This was a great moment in socializing the next generation of theoretical economists. But what *I* took

away from this experience is that I wanted to be sure that when I borrowed ideas from economics, I would take seriously how real people behaved.

Another interest of mine was artificial intelligence. This led me to read about computer chess tournaments. This in turn led to the idea that a good way to evaluate alternative strategies for the Prisoner's Dilemma would be to invite experts to submit their strategies in the form of computer programs. Then I could run a computer tournament to see which one would do best.

The result was that the simplest of all submitted entries won the tournament. This was TIT FOR TAT, the strategy that starts by cooperating, and then cooperates or defects exactly as the other player did on the preceding move. I next organized a larger tournament with both experts and computer hobbyists, with a total of 62 entries. The result was again a victory for TIT FOR TAT (Axelrod 1984).

To explore the meaning of this result, I imported some ideas from yet another discipline, evolutionary biology. The perspective of evolutionary biology allowed me to formulate and answer three questions about the how things might develop over time. The first question was how a small cluster of TIT FOR TAT players could even gain a foothold in a population of selfish players. The next question was about the conditions that would allow players using the TAT FOR TAT strategy to increase their numbers until they dominated the population. And the third question was whether such a population TIT FOR TAT players would be able to resist invasion by any player using a more selfish strategy than the simple reciprocity embodied in TIT FOR TAT. With a little high school math, I was able to answer all three of these questions, and thereby account for the emergence, growth and maintenance of reciprocity.

To extend this analysis, I once again turned to artificial intelligence, this time importing a technique for solving hard problems called the Genetic Algorithm (Holland 1975). This allowed me to confirm that the success of reciprocity was not very dependent on the particular strategies entered in the tournaments, but was likely to emerge under a very wide range of initial conditions (Axelrod 1997).

In addition to drawing on the disciplines of game theory, artificial intelligence and evolutionary biology, I found an excellent illustration of my ideas in a sociology journal. Starting in high school, I developed the habit about once a year of looking at a wide range of journals just to see what might interest me. As it happened, while thinking about what to do with my academic work on cooperation, I came across a book review in a sociology journal of a book based on the diaries and memoirs of British and German soldiers facing each other in the trench warfare World War I. What the author found was that these soldiers often employed what they called the 'Live and Let Live System.' The idea was that between large battles ordered

by the higher-ups, each side largely refrained from actions that would harm the other side – but only as long as the other side reciprocated such restraint. When I read that book, I realized that sociology provided a striking example of how opposing military units facing each other for an extended period of time could use what amounted to a TIT FOR TAT strategy to sustain cooperation even in the midst of a brutal war (Ashworth 1980).

I thought of the trench warfare example as a good way to illustrate what I was talking about. I didn't expect it would have much value for convincing people that I was on to something. After all, the trench warfare case was an historical event that I picked out of a huge number of possible historical events. But as it turned out, the trench warfare example turned out for many readers to be the most convincing evidence of all that the theory had real traction.

Having talked about how much I borrowed from different disciplines, let me turn now to what the work on the evolution of cooperation had to offer to other disciplines.

Frequently Asked Questions

But first I want to answer the three questions I am frequently asked about my work on the evolution of cooperation. These are questions about how I got the idea for the tournament in the first place, what surprised me in the course of my work on cooperation, and what I would do differently if I had it to do all over again?

The first question is where did the idea for a computer tournament come from?¹ In retrospect, I realize that it came from my interest in artificial intelligence, which started while I was in high school and an interest in game theory that started in college. In high school I came across an article about a checker-playing program that learned to improve its own play (Samuel 1959). I was fascinated. Afterwards, I followed the development of computer chess programs through the 1960s, as well as the computer chess tournaments that began in the 1970.

The second question I am frequently asked is what surprised me about my findings? Well, the first thing is that the most effective strategy was the simplest of all those submitted to the tournaments. Based upon the experience of computer chess tournaments, I had originally thought that the most effective strategy for the iterated Prisoner's Dilemma would probably be quite complex. But the opposite was true. It turned out that the very simple reciprocity of TIT FOR TAT tended to evoke cooperation from a wide range of other strategies, and therefore did well for itself.

In analyzing the millions of choices in the tournaments, I surprised to find that although TIT FOR TAT won both rounds of the tournament, it never

once did better than any of the wide range of strategies it interacted with. In fact, I came to realize that it couldn't do better than the player it was interacting with. The reason is that to do better than the other guy you are interacting with, you have to defect more than the other side does. And TIT FOR TAT never does that. It starts with cooperation and never defects unless the other guy has just defected. So it can't defect more times than the other player, and therefore can't do better than the other player. But it won both rounds of the tournament, out performing every other strategy! TIT FOR TAT is effective because it elicits more cooperation from other players than any other strategy did. Or as one reviewer said, 'Nice guys finish first.' This can't happen in a zero-sum game like chess or football, but it can and did happen in the iterated Prisoner's Dilemma. In the long run, the biggest surprise for me is how many other research programs the *Evolution of Cooperation* has inspired – something I deeply appreciate but certainly didn't expect to happen anywhere near the extent that it did. I certainly didn't expect to be standing here one day. What an honor, and what a pleasant surprise.

Adding Generosity to TIT FOR TAT

The third question I have often been asked is what I would do differently if I had to do all over again? The question is often posed something like this: 'After thirty years and tens of thousands of citations, do you see any flaws in your book on the evolution of cooperation?' The question often comes in an even more pointed form, namely, 'Do you still recommend TIT FOR TAT?' The answer to both variants of the question is that the simple reciprocity of TIT FOR TAT does need some modification in a world where misunderstanding is possible. TIT FOR TAT needs to be modified to take into account the possibility that one's cooperation will be not always will not always be perceived that way by the other player. The problem is serious because when one player thinks the other just defected, the victim is likely to retaliate, and this can lead to a series of defections that can echo back and forth for a long time.

In my design of the computer tournaments there was no possibility of random misunderstanding. Yet, some degree of misunderstanding is typical of most strategic interactions, and this can cause real trouble for TIT FOR TAT.

I wish I had earlier appreciated this basic flaw in simple reciprocity. I did not entirely ignore the problem. I did have an intuitive understanding of the problem that an occasional misunderstanding can cause. In fact, when I developed the policy implications in the book, I explicitly stated that 'in many circumstances the stability of cooperation would be enhanced if the response were slightly less than the provocation' (see Axelrod 1984, 187).

The simplest way to include the possibility of misunderstanding in the iterated Prisoner's Dilemma is to allow some percentage – say one percent or ten percent – of defections by the other player to go unpunished. This is the strategy called Generous TIT FOR TAT. Working with Ji Wu, a post-doctoral fellow from China, we showed if the original tournament had included some misunderstanding, the generous variant of TIT FOR TAT would have won the tournament anyway. It turned out that a little generosity is an effective way of dealing with the possibility of misunderstandings (see Molander 1985; Bendor *et al.* 1991; Godfray 1992; Nowak & Sigmund 1992). It also turned out that the generous variant of TIT FOR TAT was far more effective than an alternative strategy called PAVLOV that has been proposed for dealing with this problem.^{2,3}

Thus, in the presence of some occasional misunderstanding, reciprocity still works well provided it is accompanied by some generosity – which means you have some chance of cooperating when you would otherwise defect.⁴ I wish I had understood that from the beginning.

Having answered the questions I am most frequently asked about my work, I want to turn now to how the research on the evolution of cooperation has contributed to disciplines beyond political science. I'll use three examples: economics, evolutionary biology and medicine.

Contributing to Economics

Let's start with economics. Over the last three decades, many economists have become quite open to the idea that there are alternatives to the purely rational actor that has been so influential in economic theory, especially in microeconomics. Psychology offered alternatives, but these had limited appeal to economists largely because there seemed to be so many alternative psychological theories of decision making that it was hard to justify the choice of any one of them.

What has captured the imagination of many economists is the idea that economics can be studied as an evolutionary process involving a diverse population of actors who interact with each other – with results that can be studied using computational methods such as agent-based modeling. This has opened up whole new fields of evolutionary economics and evolutionary game theory. I don't want to claim any special credit here except to say that I was an early contributor to evolutionary approaches to game theory, especially in terms of broadening the scope of mainstream economics beyond its traditional fixation on so-called equilibrium states where things settle down and no longer change. I hope my exploration of how cooperation might have evolved – and may still be evolving through the development of norms and institutions (see Axelrod 1986) – has contributed a little bit to the mainstream acceptance of a much less static picture of the economy.⁵

Contributing to Evolutionary Biology⁶

Turning to contributions to evolutionary biology, I realized early on that my themes about how cooperation can evolve might have some interesting applications in biology. But as a political scientist, I also realized that I needed a lot of help to draw out these implications in a way that would be credible, and perhaps even useful to biologists. Fortunately, one of the world's leading evolutionary biologists, William Hamilton, was at the University of Michigan.⁷ So I gave him a call.

In his memoirs, Bill describes his reactions to this phone call out of the blue.⁸ And I quote:

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 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? (p. 122)

That first phone call led to a lunch where Bill Hamilton proposed that we work together. He later wrote,

Soon after the lunch . . . 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 fundamentals of whatever we were studying.

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.

Bill's disciplinary training as an evolutionary biologist and a naturalist proved essential to making our theoretical work compelling to biologists.

I shared Bill's surprise at how well we worked together. His memoirs say, and I quote:

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)

Our work in evolutionary biology has been very widely cited, and has led to empirical work finding cooperation in repeated Prisoner's Dilemma-like settings all the way from bacteria to vampire bats to primates. Our work has also led to a large number of theoretical studies by biologists that expanded on the basic model of the two-person iterated Prisoner's Dilemma. These

developments have provided new insights into other mechanisms in addition to reciprocity and kinship – such as territoriality and indirect reciprocity – that can support the emergence of cooperation within a Darwinian framework (Nowak 2006).

Contributing to Medicine⁹

It's true that I am like a little kid with a hammer who sees nails to pound everywhere I look. But I never expected to find anything 'hammerable' in the field of medicine.

The story goes like this. About ten years ago, I visited a former grad student, Stephanie Forrest, who is now one of the world's leading computer scientists. We caught up with each other, and talked about common interests, such as protecting people's privacy.

As an aside, she asked if I would be interested in seeing a computer simulation of tumor growth that she was working on. A student of hers had developed a very nice three-dimensional visual display so that you could see how a tumor mass developed over time as cells divided and mutated. It was fascinating to watch the simulated blood vessels being recruited by the tumor cells to grow in their direction. As the blood vessels grew closer, the tumor cells were able to get more than their fair share of the oxygen and nutrients that all cells need to grow and divide.¹⁰

Having worked on computer simulations of human societies, I saw what a wonderful toy Steph had. I asked what assumptions were built into the simulation, and she told me it was based on a widely cited paper called 'The Hallmarks of Cancer' (see Hanahan & Weinberg 2000).

When I went home I took a look at the paper. I was clearly in *way* over my head. But I did get the general idea, which is that cancer results from an accumulation of mutations in a single cell line.¹¹ These mutant cells achieve new capabilities, until eventually a *single* cell type is able to reproduce completely out of control.

From my experience with the study of cooperation, I had the feeling that it didn't have to be that way. I couldn't articulate what it was, but I was curious to see if there might be some fruitful analogy to be made. Playing around with it for a few weeks gave me a sense of what the analogy might be.¹² It goes like this: Any multi-cellular organism, such as ourselves, is a bit like a society in which the various parts work pretty well with each other. But in any society, there are liable to be some anti-social types who steal rather than work for an honest living. In our society, we have several ways to protect against these robbers, such as locks and alarms. My thought was that cancer might operate like a gang of specialized robbers that overcomes the various defenses society erects to protect itself.

The idea was that no single tumor cell line had to be able to overcome all the body's defenses either. Perhaps this is what accounts for how hard tumors are to control, and might even suggest a new approach to cancer therapy.

My brother Dave happens to be a cancer researcher. The next time I saw Dave I tried out my nebulous idea.¹³ Dave didn't laugh at me. And I'm grateful for that. Instead, he said my idea was not totally crazy. He then took the time to search the literature, and found that nobody had looked at cancer quite that way before.

With that encouragement, and with Dave's help, I decided to learn some of the basics about cancer, and the specialized language that cancer researchers use.¹⁴ I found that most cancer research is focused on the precise role of thousands of specific kinds of molecules. But as the details accumulate, some researchers are starting to express a real need for some additional organizing principles that could help make sense out of all the details. Perhaps Dave and I could make some progress at that level. And perhaps I could export a little social science to medicine.

Since Dave was pretty busy, I approached an oncologist at my own university, Dr. Kenneth Pienta, someone I'd heard was receptive to new theoretical approaches to cancer. Like Dave, he didn't laugh when I told him my idea about cooperation among tumor cells.

It didn't take Ken long to understand that game theory could be applied at the cellular level. The idea is that when a cell produces diffusible products for its own needs, it automatically helps nearby cells as well. So if two different cell types can each overcome different defenses in their common neighborhood, they might *both* be able to propagate faster than cells that don't cooperate. For example, one cell type might be able to overcome the control on how much blood is supplied to its neighborhood, and a nearby cell type might be able to elicit more than the normal amounts of a specific growth factor. In that case, both cell types could propagate faster than either could alone.

After a lot of hard work, we were able to show that the cooperation hypothesis is consistent with known facts about cancer; helps explain some things that hadn't been understood before, and suggests predictions that could be tested. We also showed that our approach had potential relevance to cancer therapy by suggesting the possibility of interrupting the cooperation between tumor cells.

We showed our paper to a few other cancer researchers and got very positive feedback. They said in effect, 'Hey, it's obvious *after* you say it.' But when we submitted it for publication, we got a flat rejection. I've gotten my share of rejections, but every time it seems like a kick in the stomach. In this case, one of the reviewers said: everyone knew this already. And the other reviewer said it was absolutely impossible.

Well, after we picked ourselves up off the floor, we revised the paper to make clear exactly what was new, and why it wasn't impossible. The revised version came out in another journal (Axelrod *et al.* 2006).

In the last few years, according to Dr. Pienta, 'The theory of cooperation within the cancer ecosystem has gained tremendous traction. Experiments have . . . proven that cooperation exists between cancer cells as well as [between] cancer cells and host cells to facilitate tumor growth. The work is causing a paradigm shift in thinking about [how tumors get started].'

Additional Prospects for Political Science Contributions

Having described how I have used political science ideas to contribute to the well-established disciplines of economics, evolutionary biology, and medicine, let me suggest that political science has a lot to offer new fields as well.

- For example, neuropsychology is beginning to understand how, in a laboratory, a person's perception of fairness can operate at the neural level. Political science has a lot to offer for extending this new understanding from the neural to the interpersonal level.
- Social media such as Twitter and Facebook have over a billion users who together have created a huge new commons in cyber space. Political science has a lot to offer on what it takes for users themselves to govern such commons See Ostrom (1990).
- As a third example, cyber security depends in large part on the willingness of individual computer users to take the trouble to maintain the security of their own computers, and political science has a lot to offer about how to promote pro-social norms and metanorms in large societies See, for example, Axelrod (1986).

To conclude, political scientists have gained tremendously by borrowing from established disciplines, but we also have quite a lot to offer new as well as established fields that seek to understand how to better the human condition.

NOTES

1. This section is adapted from Axelrod (2009). Reprinted with permission of Elsevier Publishing. See this article for more on the origins and development of my work on the evolution of cooperation.
2. PAVLOV is also known as Simpleton (Rapoport & Chammah 1965), and Win Stay Lose Shift (Nowak & Sigmund 1993). This rule starts with cooperation. From then on, it changes its choice if and only if it received one of the two lowest payoffs (i.e., the

other player defected). This rule did well in a non-standard simulation that did not take account of discounting of payoffs over time (Nowak & Sigmund 1993). However, neither the basic nor the generous version of the strategy did well when some misunderstanding was added to the heterogeneous environment of the second round of the Prisoner's Dilemma Tournament (Wu & Axelrod 1995). In fact, with the usual payoff parameters it is easily exploited by a strategy that alternates between cooperation and defection (Rapoport & Chammah 1965).

3. If a player realizes that its intention to cooperate was for some reason not implemented, a little contrition can help. See Wu and Axelrod (1995) for another slight modification of TIT FOR TAT called Contribute TIT can be very effective.
4. If a player realizes that its intention to cooperate was for some reason not implemented, a little contrition can help. See Wu and Axelrod (1995).
5. Of course economists have long been interested in business cycles, but the standard treatments assume that these are merely *cycles* rather than progressions to new levels of organization.
6. This section is adapted from Axelrod (2006). Reprinted by permission.
7. 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).
8. All quotes are from Hamilton (2001).
9. This section is adapted from Axelrod (2008). Reprinted by permission of Cambridge University Press.
10. Each different kind of mutant cell was represented with a different color. You could rotate the image in three dimensions. Neatest of all, when you clicked on a cell, all the cells of that type would become transparent so that you could see inside the simulated tumor.
11. The literature typically does not explicitly say that a single cell needs to have all the hallmarks of cancer, but this seems to be the tacit assumption.
12. My first analogy was inspired by Hobbes' idea that a healthy society is analogous to the healthy body of an individual that he called the Leviathan. The analogy was that cooperation among healthy cells would help the body suppress tumors. But that didn't get me beyond what biologists already know, e.g., how multicellular organisms sustain cooperation among their more or less specialized cells. When that approach didn't take me to new territory, I tried another approach, namely to think about the possibility of cooperation among the mutant cells rather than among the normal cells.
13. At that point, the best I could do was to ask why if robbers don't have to go it alone, why should tumor cells? Just as robbers cooperate with each other, why couldn't tumor cells cooperate with each other?
14. Dave's first recommendation was Greaves (2000).

REFERENCES

- Ashworth, T. 1980. *Trench Warfare, 1914–18: The Live and Let Live System*. New York: Holmes & Meier.
- Axelrod, R. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, R. 1986. 'An Evolutionary Approach to Norms', *American Political Science Review* 80, 1095–111.
- Axelrod, R. 1997. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.
- Axelrod, R. 2006. 'Agent-based Modeling as a Bridge between Disciplines', in Judd, K. L. & Tesfatsion, L., eds, *Handbook of Computational Economics II: Agent-based Computational Economics*. New York: North-Holland.
- Axelrod, R. 2008. 'Political Science and Beyond: Presidential Address to the American Political Science Association', *Perspectives on Politics* 6, 3–9.

- Axelrod, R. 2009. 'Launching "The Evolution of Cooperation"', *Journal of Theoretical Biology* 299, 21–4.
- Axelrod, R., Axelrod, D. E. & Pienta, K. J. 2006. 'Evolution of Cooperation among Tumor Cells', *Proceedings of the National Academy of Sciences* 103, 13474–79.
- Bendor, J., Kramer, R. M. & Stout, S. 1991. 'When in Doubt: Cooperation in the Noisy Prisoner's Dilemma', *Journal of Conflict Resolution* 35, 691–719.
- Godfray, H. C. J. 1992. 'The Evolution of Forgiveness', *Nature* 355, 206–7.
- Greaves, M. 2000. *Cancer: The Evolutionary Legacy*. Oxford: Oxford University Press.
- Hamilton, W. D. 1964. 'The Genetical Evolution of Social Behavior', *Journal of Theoretical Biology* 7, 1–52.
- Hamilton, W. D. 2001. *Narrow Roads of Gene Land: The Collected Papers of W. D. Hamilton, Vol. 2: The Evolution of Sex*. Oxford: Oxford University Press.
- Hanahan, D. & Weinberg, R. A. 2000. 'The Hallmarks of Cancer', *Cell* 100, 57–70.
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.
- Molander, P. 1985. 'The Optimal Level of Generosity in a Selfish, Uncertain Environment', *Journal of Conflict Resolution* 29, 611–8.
- Nowak, M. 2006. 'Five Rules for the Evolution of Cooperation', *Science* 314, 1560–63.
- Nowak, M. & Sigmund, S. 1992. 'Tit for Tat in Heterogeneous Populations', *Nature* 355, 250–3.
- Nowak, M. & Sigmund, S. 1993. 'A Strategy of Win–Stay, Lose–Shift that Outperforms Tit-for-Tat in the Prisoner's Dilemma Game', *Nature* 364, 56–8.
- Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- Rapoport, A. & Chammah, A. M. 1965. *Prisoner's Dilemma: A Study in Conflict and Cooperation*. Ann Arbor, MI: University of Michigan Press.
- Samuel, A. L. 1959. 'Some Studies in Machine Learning Using the Game of Checkers, II: Recent Progress', *Annual Review in Automatic Programming* 6 (Part 1), 1–36.
- Trivers, R. 2000. 'Obituary: William Donald Hamilton (1936–2000)', *Nature* 404, 828.
- Wu, J. & Axelrod, R. 1995. 'How to Cope with Noise in the Iterated Prisoner's Dilemma', *Journal of Conflict Resolution* 39, 183–9.