Understanding the Limits of Competitive Processes

Henry Lieberman
Christopher Fry
MIT Media Lab
Cambridge, MA 02139 USA
{lieber, cfry}@media.mit.edu

ABSTRACT

Darwin was half-right. He stressed the importance of competitive processes in evolution, but he may not have fully appreciated the complementary role that cooperative processes play, from the emergence of multicellular organisms, to symbiosis.

Today, our economics and politics laud competitive processes. They tell us that if each person or organization does what is in their own interests (and everybody "competes"), the result will be globally best. That assumption, while true in some circumstances, has its limits. We are increasingly bumping up against these limits, with potentially disastrous consequences for society.

Game theory gives us a perspective for understanding the tradeoffs between competition and cooperation. But it isn't easy for people to arrive at an understanding of the economic and social consequences, because it involves shifting from a local perspective to a more global perspective.

Computation can help, because it allows us to "jump out of our skin" enough to see general principles. Computation helps us manage the task of determining these tradeoffs, and applying the principles in practice. We can

- Simulate the consequences of various cooperation vs. competition tradeoffs, and mathematically analyze them;
- Use online decision support systems to discuss tradeoffs, get the perspective of others, and collaborate on solutions;
- Develop cooperative alternatives to formerly competitive processes;
- Educate people about the underlying principles and their real-world consequences.

INTRODUCTION

Contemporary western society idolizes competition. A common meme is that "competition brings out the best". Our capitalist economy is based on competition between products and competition between companies. Our political system of representative democracy is based on competition between political parties and competition between candidates.

While there is a lot of truth to the advantages of competition, there are limits. The worst thing about competition is that it "competes against" cooperation – and cooperation can lead to win-win outcomes. A more nuanced view of competition, understanding its limits, can lead to wide-ranging improvements in political and economic systems.

The game theory concept of the *Iterated Prisoners' Dilemma* (IPD) provides us a framework for understanding the limits of competitive processes. We start by talking about a related economic hypothesis we call the Fundamental Theorem of Capitalism (FToC), and explore its consequences. We then discuss how modern computation and communications technologies can change the balance between competition and cooperation, for the better.

	He cooperates	He doesn't cooperate
I cooperate	*Reward:* We both get **$3** ! :-)	*Sucker's payoff:* I get **$0** He gets **$5**
I don't cooperate	*Temptation:* I get **$5**, He gets **$0**	*Punishment:* We both get **$1** :-(

Figure 1. The Prisoners' Dilemma. Two players are given a choice whether to cooperate or not. I notice that I am better off *not* cooperating regardless of whether the other player cooperates or not ($5>$3, $1>$0). But he does the same thing, the result is we both get $1. Whereas if we both cooperated, we could have gotten $3 each!

THE "FUNDAMENTAL THEOREM OF CAPITALISM"

Laissez-faire capitalism and contemporary neoliberalism are based on an assumption -- that if each person or organization does what is narrowly in their best economic interests, the result will be globally optimal. Let's call this the Fundamental Theorem of Capitalism (FToC). We now know that, under many conditions, the FToC is false.

Modern game theory, especially the Iterated Prisoners' Dilemma (IPD) [Axelrod 1984], teaches us that (providing certain inequalities hold), if each agent does what is locally optimal for themselves, the result can be *worse* for everyone than if everyone cooperated! (In the IPD, competition (or failure to cooperate) is technically referred to as *defection*).

Figure 1 shows a single instance of a cooperation-vs-competition game. The situation we're talking about occurs whenever

Temptation > Reward > Punishment > Sucker's payoff

The far more interesting case is when the game is *iterated*, and players must choose a long-term strategy.

This is a mathematical result. It is not a political position and it is not debatable. (Though what the inequalities actually amount to in a given real situation is certainly open for debate). It constitutes a fundamental limitation on what competitive processes such as markets can achieve. Today's capitalist markets and adversarial political structures such as elections or court cases simply operate as if no such limits exist.

We maintain that blind adherence to the FToC is emblematic of unsustainable social, economic, and ecological practices, such as war, pollution, racial and other discrimination, destructive commercial competition, income inequality, and a host of other societal ills. We list a few.

• *Commercial competition*: It is often thought that competition "causes the best products to win". Sometimes it does. But this kind of *economic Darwinism* also causes duplication of effort, misleading and false advertising, exploitation of consumers, etc. Nobody knows whether the purported advantages of competition outweigh these disadvantages, because this tradeoff is never taken into consideration.

Economic Darwinism (and any form of Social Darwinism) is based on an antiquated view of evolutionary theory. Modern evolutionary theory says evolution selects, not just for (the circular criteria) "survival of the fittest", but for positive-sum games [Wright 01].

• *War*: If each nation computes the consequences of being attacked versus the cost of military preparation, it opts for military preparation. But that preparation itself increases the likelihood of a war (a military-industrial complex advocating war, demonization of foreign groups, exaggeration of threats, etc.), causing both sides to lose. Sustainable peace would free resources for both sides.

• *Pollution*: A polluter gets benefit from the activity causing the pollution, such as industrial production, while everyone else bears the cost. But if everyone pollutes, global warming or other ecological disruption causes everyone to lose, including the polluter. These kinds of situations are what economists call *externalities*.

• *Adversarial political processes*: Witness the gridlock in today's US Congress. Republicans and Democrats can't agree on anything any more than the Yankees and the Red Sox can agree on who should win a baseball game. Politicians see their main job as to win zero-sum elections rather than to collaborate on solving the country's problems.

It is often to the short-term advantage of an individual politician or lawyer to pander, mislead, evade, (even lie) or sell themselves to special interests. But if it becomes a social norm in political discourse, the poor voter or juror who has to make a decision is faced with a game of "liar's poker", and has no basis for making a rational decision.

• *Racism*: Even phenomena like racism can be modeled with the IPD. [Axelrod & Hammond 03] reported an IPD simulation where a majority group cooperated with each other and defected with the minority group. It has nothing to do with the actual characteristics of the group, only that it is apparent which group a given agent belongs to. It formed a stable pattern. But of course, it led to suboptimal results for everybody. Similarly any kind of nationalism, tribalism, sexism, homophobia, religious discrimination, is a loss for everybody.

Above, we have been emphasizing the negative consequences of competition. To be fair, we should note that there's a flip side, where cooperation fails and competition succeeds. Again, it just depends on what the numbers are in the Prisoners' Dilemma matrix. Libertarians are quick to point out things like Hardin's *Tragedy of the Commons*, [Hardin 68] which is actually a corollary of the IPD. They use this to explain, for example, the failure of 20th century Communism, a positive case for the FToC.

Wright [Wright 01] makes the case, on evolutionary grounds, that the tradeoff is determined by the availability of resources – competition is best if resources are scarce, cooperation if resources are abundant. We agree, and believe it is the abundance of (computational) resources in the 21st century that will tip the scales in favor of cooperation.

It is liberating to realize that all these problems are instances of a common pattern. So much of today's political and economic discourse accuses the opposition of malicious intent or moral failings. Actually, it's nobody's fault. It's just that we're stuck in a pattern that we haven't been able to see. Once we do see it, we can shift the conversation from non-issues like, "Is cooperation or competition better?" in general, to trying to understand the tradeoffs, and managing them in a positive and sustainable way.

DOES COMPETITION MOTIVATE?

Another half-true cultural myth is that "competition motivates people". Certainly, in some situations such as competitive games, that's true. But what kind of motivation does competition provide?

Competition doesn't motivate all people equally. It works best with people who have "competitive personalities", which have their good and bad sides: drive and determination, yes, but also aggression and hostility. Competitive personalities tend to be more associated with men rather than women. Blanket assertions that people will be motivated by competition tend to disenfranchise those who don't fit the competitive personality profile.

There's nothing wrong with competition as entertainment, when games are good clean fun between consenting adults. But in *gamification*, which artificially introduces competition in education and the workplace, people may feel obligated to participate. Those who don't have competitive personalities will actually be *demotivated* by artificially competitive situations. They sense, not incorrectly, that situations that necessarily have few winners and many losers can be a sucker bet.

Psychologists distinguish between *intrinsic* and *extrinsic* motivation [Kohn 92]. Intrinsic motivation for an activity means that you want to do the activity for its own sake. You listen to music because you enjoy hearing it. Extrinsic motivation is provided by incentives that are external to the activity itself: rewards, prizes, grades, and rankings. Competition can only provide extrinsic motivation.

Another way to view the results of the IPD is as a trade-off between short-term and long-term. Defection accomplishes a short-term gain, at the expense of the long-term opportunity cost of missing out on the benefits of cooperation.

Numerous studies have shown that while extrinsic motivation can be effective in the short term, it tends to decrease intrinsic motivation in the long term. Kohn [Kohn 93] describes an experiment where young children were given a dollar for each crayon drawing they produced. Initially they produced more drawings than a control group that was not rewarded. But then, much later, simply left in a room with crayons and paper, and offered no reward, they were far less motivated to draw than the control group!

SIMULATION

We are also faced with another kind of limit – limits on our computational ability as humans to understand the complexity inherent in cooperation-competition tradeoffs. Each person has only limited time, limited knowledge, and limited ability to perform inference.

One reason why the FToC and its ilk are so seductive is that they only require that each individual agent consider its own perspective, not that of others, which makes it a computationally less challenging task. Locality in general is a good thing, except when it causes you to miss an important nonlocal property, as it does here. Fortunately, though, we can use the computational power of our machines to help us meet the challenges of understanding global perspectives.

The appreciation of global perspectives often causes a sea change in people's attitudes, helping them get beyond a local bias. A watershed moment in the ecology movement was when the first Blue Marble picture (Earth from space) was released—dramatizing the finiteness of the planet and the interconnection of our shared presence upon it.

To begin with, we would not even have the understanding of the IPD we do today, were it not for computational assistance. The mathematics of the Prisoners' Dilemma was first discovered in 1950 by Merrill Flood and Melvin Dresher (ironically enough for libertarians, at the Rand Corporation!). They were demonstrating the folly of the mutually-assured-destruction nuclear war strategy. But it wasn't until Axelrod's 1984 book, which reported a series of computational experiments with simulated agents following a variety of strategies, that the game theory perspective emerged, leading to our current understanding.

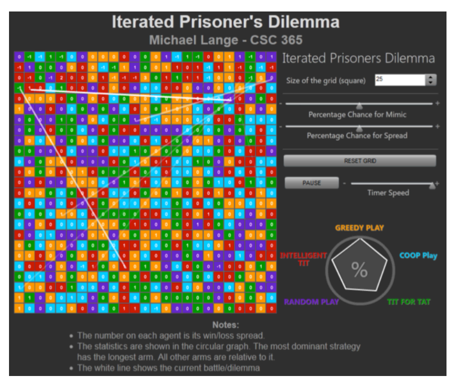

Figure 2. An interactive simulation of various strategies for the Iterated Prisoners' Dilemma.

In Axelrod's simulations, the most successful strategy was called *Tit-for-Tat*, which started out by cooperating, and then reciprocated the opponent's behavior. In general, successful strategies were *nice* (never the first to defect), *provokable* (avoided being taken advantage of), and *forgiving* (willing to try cooperating in the face of defection in the hopes of inducing others to cooperate). We should note that the theoretical problem of an optimal strategy for the IPD in general has not yet been solved, and remains a topic of active research. However, we already have more than enough knowledge (such as the success of TfT and nice/provokable/forgiving strategy attributes) to provide important practical lessons for society.

Although the basic results concerning the IPD are familiar to mathematicians, we believe that the general public (and, especially, our business and political "leaders") do not appreciate the extent to which IPD dynamics really govern many economic and political issues. We believe that computer tools based on a deep understanding of the IPD could have a profound practical effect on helping people make better decisions.

Just as a simulation of IPD agents helps us understand the mathematical problem in the abstract, we can also use simulation to understand how different strategies might play out in the real world. Should we raise or lower taxes? Rather than get into abstract moral discussions about how much individuals or businesses truly "owe" to the government, if we had a realistic economic simulation, we could try various strategies and see what the consequences are. We are not at the point today where we can make economic simulations that are realistic and understandable enough to be able to answer these kinds of questions definitively. But it is better to center a debate about the fidelity, assumptions, and outcomes of simulations, than about some preconceived ideological notion of the proper size of government.

DECISION SUPPORT

Computation can also help us in bridging the gap between abstract mathematical understanding and real-world political and economic discourse. Every day, people have to make decisions about whether to cooperate or to compete, or whether to support cooperation or competition in their organizations. They need to discuss these decisions, and come to an understanding with others. This involves a lot of complexity, and we believe interactive decision-support systems can be a vital tool in coping with this complexity.

The primary decision procedure today in most organizations is the so-called "meeting". This involves a relatively unstructured real-time discourse, where everybody gets to express their opinion, and then a decision is made. After such a meeting, people often have difficulty remembering who said what, what was addressed, and why decisions were made. Usually, the decision procedure is either the democratic *emote and vote*, or, if there's a designated authority, *plea and decree*. Neither provides the best opportunity for creatively solving problems and achieving consensus [Susskind and Cruikshank 06].

Decision support tools can help with recording rationale in a structured fashion, help people who join mid-discussion to catch up, and decouple decisions from the pressure of real-time response, personalities and emotions. As with simulations, decision support tools can help with exploring the consequences of what-if scenarios and cooperation-competition tradeoffs.

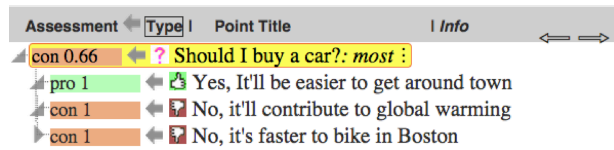

Figure 3. The Justify decision support system. Here, an argument weighing the pros and cons of buying a car. Since there are two reasons against and a single reason in favor, Justify summarizes the decision as "no", with a weight of 2/3.

Justify [Fry and Lieberman 13] is a decision support system that records a discussion as a hierarchy of *points*, each expressing a single question, fact, or opinion. Each point has a *type* that expresses its role in the discussion, such as *pro* or *con* its superpoint in the hierarchy. Justify provides semi-automatic summarization at every level of the hierarchy, and, like a programming language, an interactive development environment (IDE) for a variety of decision procedures.

One of our research goals for Justify is to provide explicit support for discussion and negotiation patterns that are more likely to lead to productive, "win-win" cooperation. Adversarial discussions where each side tries to win at the expense of the other are rarely convincing, and doomed to get decided only by "might makes right" power relationships ("I've got more votes than you" is also a power relationship).

The Harvard Project on Negotiation (www.pon.harvard.edu) and Consensus Building Institute (www.cbuilding.org) have a long history of work in translating the lessons of game theory for a general audience. They offer practical advice to business and political leaders aimed at encouraging win-win cooperation and defusing adversarial interactions. PoN is best known for its series of popular books starting with "Getting to Yes" [Fisher, Ury, Patton 91]. Other threads of work come to similar conclusions from the alternative perspectives of counterculture politics [Butler and Rothstein 87] and psychotherapy [Rosenberg 03]. These books provide sets of guidelines to be followed in meetings and other person-to-person communication. We are investigating whether the kinds of communication patterns recommended in these books can be given explicit computational support in systems like Justify.

For example, one issue concerns the idea of *maximizing* goals in negotiations. Conventional negotiation theory assumes that each party should want to get as much as possible for themselves out of the negotiation – "More is better".

But this is another case, like the FToC, where natural limits to the process are not being recognized. In fact, more is not always better [Schwartz et al 02]. Most true "utility curves" are not completely linear nor infinite. More commonly, there's a certain level, below which a proposal is unacceptable, the minimum needed to *satisfice* the goal. Beyond that, there's a linear range where, indeed, more *is* better. But, beyond a certain level, it plateaus. Holding out for more beyond the plateau only serves to make agreement more difficult.

In negotiating salary, for example, employees are expected to try to get as much money as possible. However, surveys show that life satisfaction plateaus at a relatively modest level of income – around $50-75K/year [Kahneman and Deaton 10]. Conversely, some employers support the idea of a minimum wage or living wage for their employees, despite having enough market power to force employees with few alternatives to accept lower salaries.

The problem with maximizing goals is that it sounds good when you hear it from the football coach, but if all agents do it, it forces a zero-sum game, risking mutual defection as in the IPD. One possible role for a computer agent (as is sometimes the case for a human mediator) is to elicit the utility curves from each party independently (so they don't have to disclose them to each other, which, in an adversarial negotiation, would cede market power). Then the mediating agent could compute the "trading zone" where all parties have their needs satisficed, and the surplus value is fairly shared.

Susskind [Susskind 14] and others advocate multi-attribute negotiations, where win-win outcomes can arise from differences between how the parties value the different attributes ("You get the flesh of the lemon for juice, and I get the peel for my cake"). But multiple attributes and the combinatorics of attribute bundles increase the potential complexity of negotiations, which again, cries out for computer support.

COOPERATIVE ENTERPRISE

Once we understand that there is substantial benefit in exploiting untapped opportunities for cooperation, how do we put it into practice, especially when the dominant economy assumes a competitive stance? The danger is that small numbers of cooperators can be defeated by defection from a larger group of competitors, a danger that is often borne out by observation of IPD simulations.

One long-standing answer is the formation of *cooperatives*. I (Lieberman) have my bank account and my mortgage in a cooperative bank (credit union), buy my food from a food co-op, shop at a cooperative university bookstore, have lived in housing cooperatives for decades, have bought car insurance from a cooperative, and have my bike fixed at a cooperative repair shop (which also offers to teach me how to fix it myself).

Fortunately, our society permits nonprofit cooperatives, though they still face discrimination from conventional economic institutions. Cooperatives provide viable alternatives to competitive economic institutions. They can be started on a small scale and grown incrementally. They don't require violent revolutions or mass protest movements, and can co-exist with capitalist institutions. The Internet itself is perhaps the best and most impactful example of a cooperative. It took over from for-profit, competing information services (anyone remember Compuserve and The Source?).

Cooperation requires coordination. Coordination is work, so sometimes the reason we don't cooperate is simply that it's more work. As the cost of communication falls, this decreases the cost of cooperation, making it more attractive. In the pre-internet, pre-cell phone economy, the expense of communication was somewhat mitigated by centralizing coordination, usually by a company supplying that coordination as a service for profit. This is not bad in itself, but the problem is that if coordination is centralized in one (or a few) companies, they will have a tremendous temptation to use their oligopoly position to extract an inordinate cost from their customers.

Game theory models this through what is called the *Ultimatum Game* [Guth et al 82]. That is, the company will compete with its customers, because it can, or is pressured to by Wall Street. If there's only one company that supplies phone service to your house, you'll likely pay a high cost to the company even though that's unfair. As we have seen, this results in a situation that is worse for everybody.

The so-called "New Economy" enabled by the Internet refers to the fact that coordination can increasingly be provided by distributed computing. We can cut out the very expensive middlemen, called *disintermediation*.

Examples abound. Travel agents were disintermediated by airline reservation sites. Amazon disintermediated bookstores. YouTube disintermediated television. AirBnB disintermediated hotels. Zipcar, Uber, and bike share disintermediated transportation. The evolution of 3D printers may disintermediate most manufacturing companies.

Disintermediation represents an enormous opportunity, but there are pitfalls. The new for-profit intermediaries will have to resist the temptation (predicted by game theory) to become monopolists themselves (are you listening, Amazon?). And disintermediation, like any form of automation in a capitalist society, may reduce the number of "jobs" and we need new mechanisms to make sure that innocent people do not suffer as a result.

In general, society would do well to encourage the formation of cooperatives and low-overhead private intermediaries. We should study such situations, and understand the principles that cause them to succeed or fail, relative to traditional capitalist profit-maximizing enterprises. We should be on the alert for unfair defensive actions by traditional companies, which will try to put obstacles in the path of new institutions (via regulation, bribery, cartels, etc.). And we should put a priority on the development of new technologies that will facilitate cooperation and disintermediation, such as automatic matching of supply and demand, and do-it-yourself products and services.

EDUCATION

Last but not least, our most powerful tool in encouraging a more cooperative society, is simply, *education*. Fortunately, that too, is being facilitated and made more accessible by technological progress.

We recommend that the fundamentals of game theory, including the Prisoner's Dilemma, be taught to everybody. The math is easily understandable to high-school students.

It can be fun to teach, starting with role-playing games, where students have to make the choice of whether or not to cooperate in a face-to-face situation. They can play with computer simulations of strategies, just like researchers do, encouraging a hands-on, experimental approach. Students can be challenged to think about how these principles apply in their own lives and in the situations they see around them and read about. Online courses such as MOOCs can also make educational materials about this subject accessible to anyone connected to the global Internet.

CONCLUSION

Our present economic and political institutions were developed at the time of the Industrial Revolution, where material resources were relatively scarce, communication and coordination were relatively difficult, and little was known about the mathematical structure of competitive and cooperative situations. Now we are in the Information Age, where information resources are abundant, communication and computation are inexpensive, and we have a much better theoretical understanding of game theory. It's time to rethink economics and politics.

Society needs both competitive and cooperative processes. But we are now in an age of transition, where the balance between competition and cooperation is changing. We no longer have to settle for the limits of a society that always prioritizes competition over cooperation. We believe that, in no small part as a result of advances in science and technology, the balance is shifting radically in favor of cooperation. If we all win, each of us wins.

REFERENCES

1. Robert Axelrod, *The Evolution of Cooperation*, Basic Books, 1984.
2. Robert Axelrod and Ross Hammond, The Evolution of Ethnocentric Behavior, Midwest Political Science Convention, 2003.
3. C.T. Butler and Amy Rothstein, *On Conflict and Consensus*, Food Not Bombs Publishing, 1987.
4. Roger Fisher, William Ury, Bruce Patton, *Getting to Yes: Negotiating Agreement Without Giving In*, Penguin, 1991.
5. Christopher Fry and Henry Lieberman, Decision-Making Should be More Like Programming, Symposium on End-User Programming, Copenhagen, 2013.
6. W. Güth, Schmittberger, and Schwarze (1982). "An Experimental Analysis of Ultimatum Bargaining". Journal of Economic Behavior and Organization 3 (4): 367–388.
7. Garrett Hardin, The Tragedy of the Commons, Science (AAAS) 162 (3859): 1243–1248., 1968.
8. Daniel Kahneman and Angus Deaton, High income improves evaluation of life but not emotional well-being, Proceedings of the National Academy of Sciences, Vol. 107, No. 38, 2010.
9. Alfie Kohn, *No Contest: The Case Against Competition*. Houghton Mifflin, 1992.
10. Alfie Kohn, *Punished by Rewards: The Trouble with Gold Stars, Incentive Plans, A's, Praise, and Other Bribes*, Houghton Mifflin, 1993.
11. Marshall Rosenberg, *Non-Violent Communication*, Puddle Dancer Press, 2003.
12. B. Schwartz, A. Ward, J. Monterroso, S. Lyubormirsky, K. White, D. Lehnman, Maximizing vs. Satificing: Happiness is a Matter of Choice, Journal of Personality and Social Psychology, Vol. 83, No. 5, pp. 1178-1197, 2002.
13. Lawrence Susskind and Jeffrey Cruikshank, *Breaking Robert's Rules*, Oxford 2006.
14. Lawrence Susskind, *Good for You, Great for Me*, Public Affairs Press, 2014.
15. Robert Wright, Nonzero: The Logic of Human Destiny, Knopf, 2001.

.